



US010509785B2

(12) **United States Patent**  
**Arye et al.**

(10) **Patent No.:** **US 10,509,785 B2**  
(45) **Date of Patent:** **\*Dec. 17, 2019**

(54) **POLICY-DRIVEN DATA MANIPULATION IN TIME-SERIES DATABASE SYSTEMS**

(71) Applicant: **Timescale, Inc.**, New York, NY (US)

(72) Inventors: **Matvey Arye**, New York, NY (US); **Michael J. Freedman**, Princeton, NJ (US); **Robert Kiefer**, Hoboken, NJ (US); **Ajay A. Kulkarni**, New York, NY (US); **Erik Nordström**, Stockholm (SE)

(73) Assignee: **Timescale, Inc.**, New York, NY (US)

(\*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

This patent is subject to a terminal disclaimer.

(21) Appl. No.: **16/109,644**

(22) Filed: **Aug. 22, 2018**

(65) **Prior Publication Data**  
US 2019/0188204 A1 Jun. 20, 2019

**Related U.S. Application Data**  
(63) Continuation of application No. 15/907,105, filed on Feb. 27, 2018, now Pat. No. 10,073,903.  
(Continued)

(51) **Int. Cl.**  
**G06F 16/30** (2019.01)  
**G06F 16/2455** (2019.01)  
(Continued)

(52) **U.S. Cl.**  
CPC .... **G06F 16/24554** (2019.01); **G06F 16/2264** (2019.01); **G06F 16/2272** (2019.01);  
(Continued)

(58) **Field of Classification Search**  
CPC ..... G06F 17/30486  
See application file for complete search history.

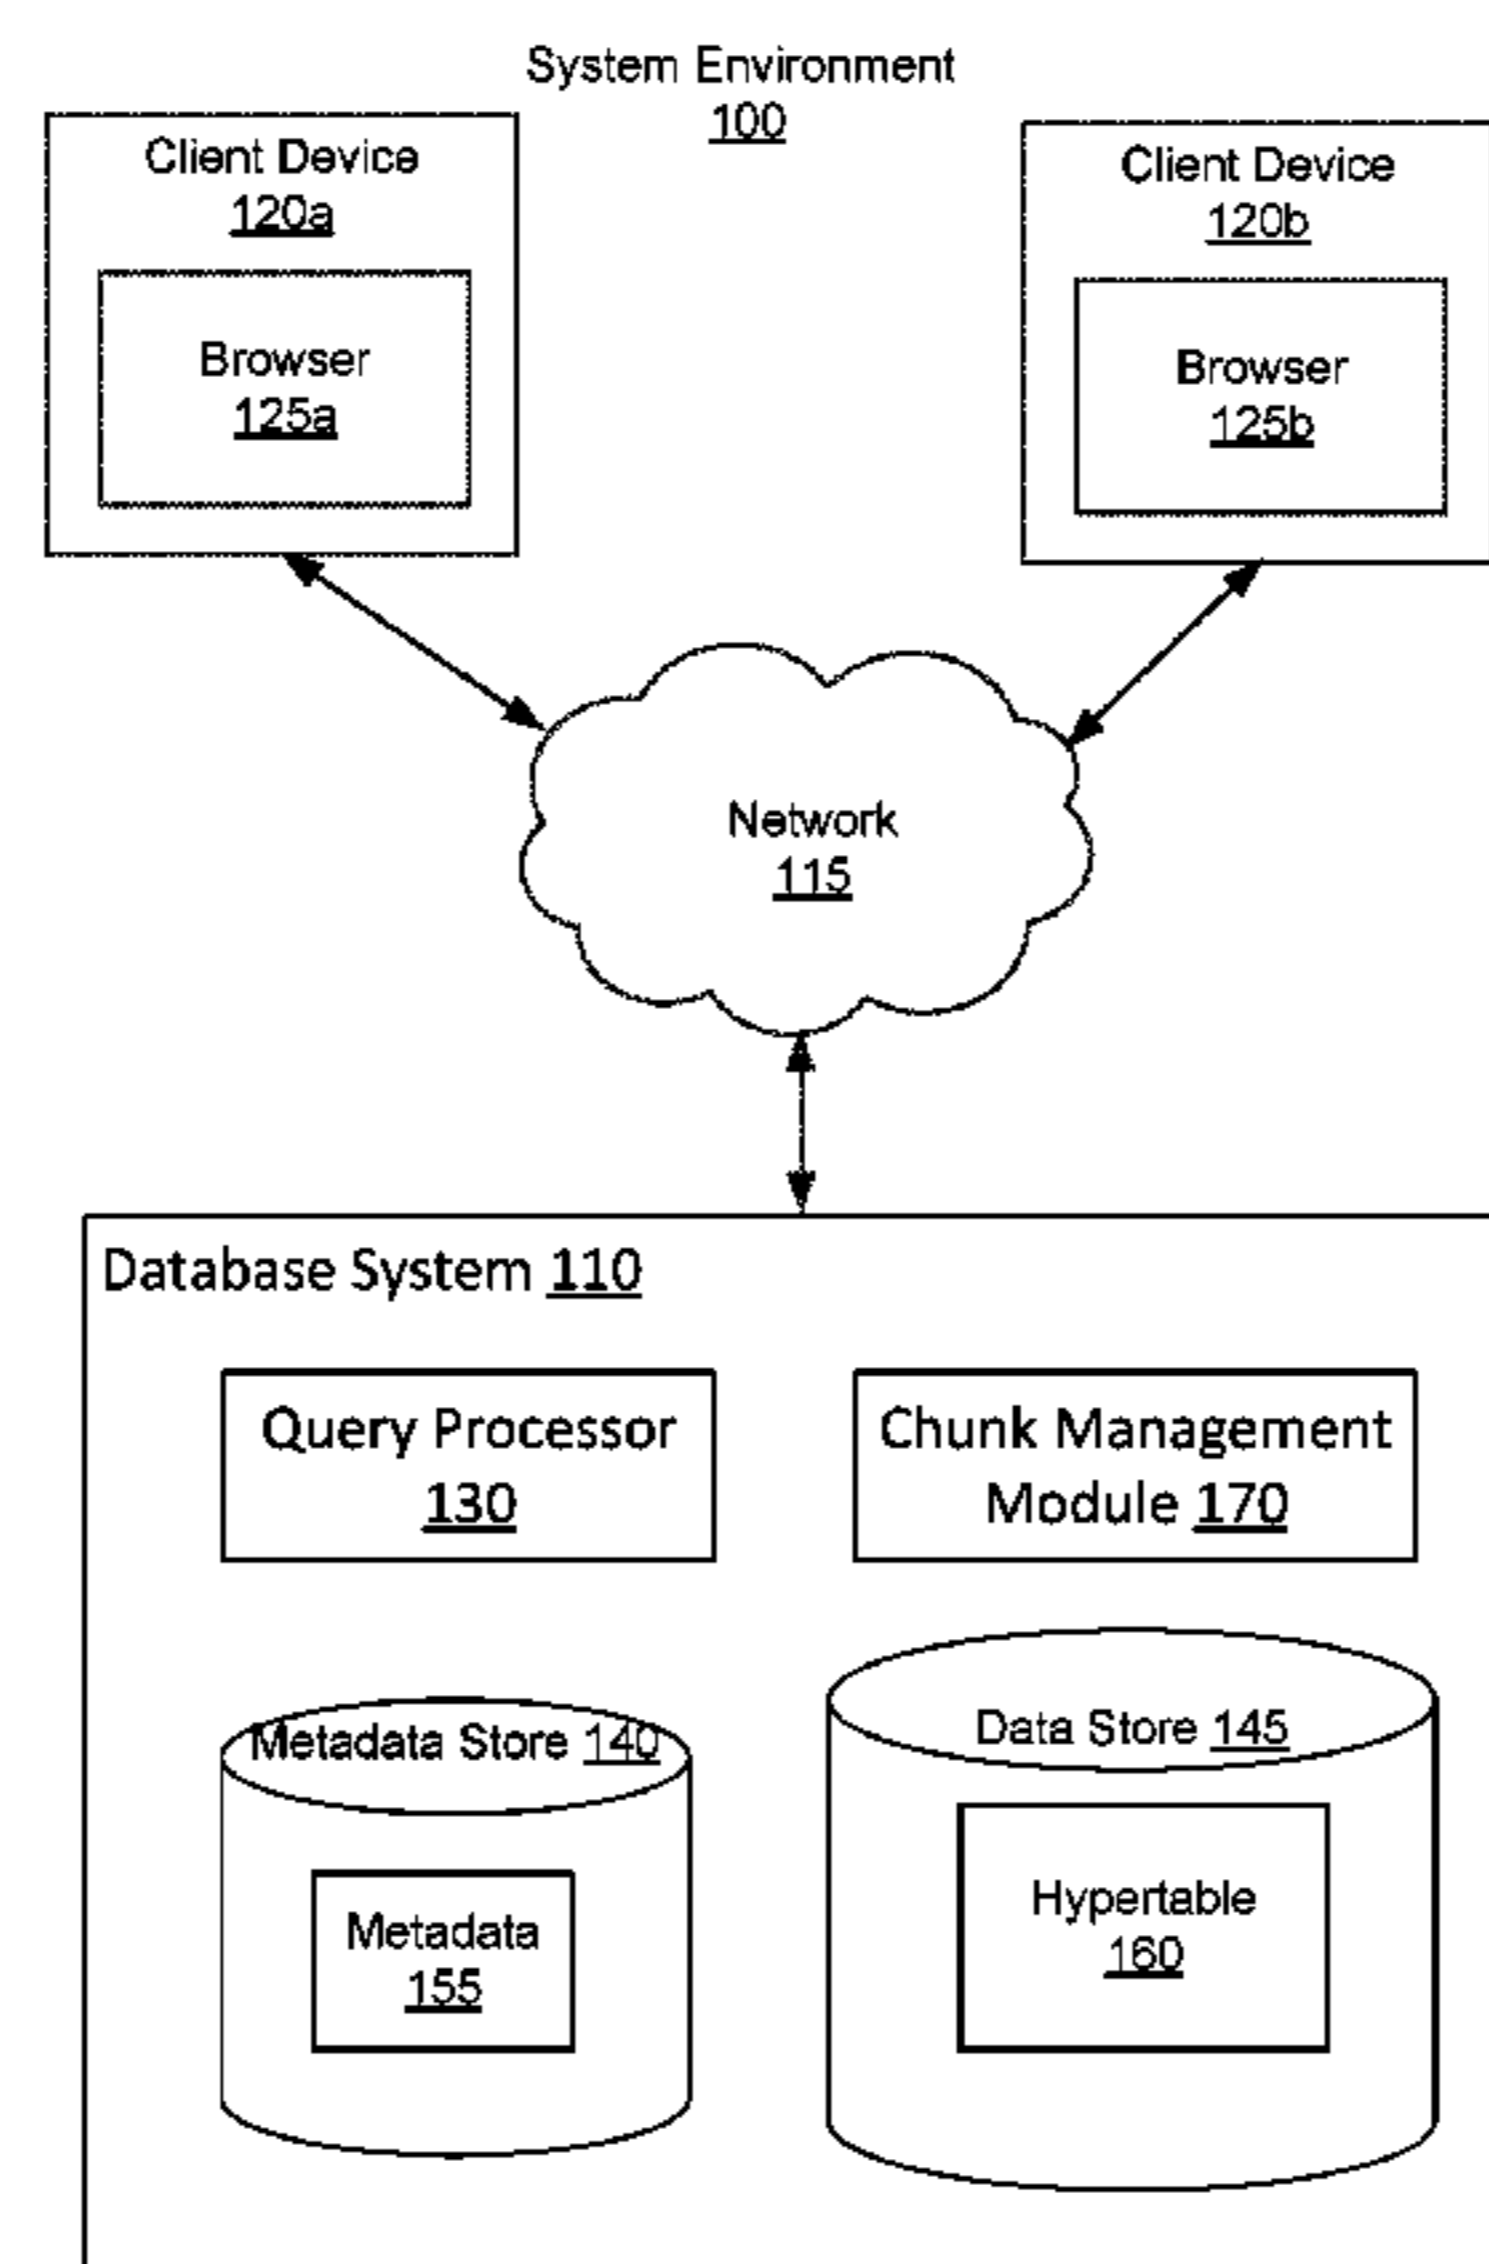
(56) **References Cited**  
U.S. PATENT DOCUMENTS  
9,740,762 B2 \* 8/2017 Horowitz ..... G06F 16/278  
2007/0198591 A1 8/2007 Teng et al.  
(Continued)

**OTHER PUBLICATIONS**  
PCT International Search Report and Written Opinion, PCT Application No. PCT/US2018/019990, dated Jun. 11, 2018, 28 pages.  
(Continued)

*Primary Examiner* — Mahesh H Dwivedi  
(74) *Attorney, Agent, or Firm* — Fenwick & West LLP

(57) **ABSTRACT**  
A database system stores data as hypertables that represent partitioned database tables. Each hypertable comprises chunks of data that may be distributed across multiple locations, each location comprising at least a storage device. The database system provides an interface that allows database queries seamlessly to hypertables as well as standard tables. The database system dynamically creates chunks as records are added to a hypertable. The database system defines a new partitioning strategy if the storage configuration of the database system is changed by adding new locations or removing existing locations. The records added to the hypertable before the storage configuration was changed continue to be stored as chunks distributed according to the previous partitioning policy.

**48 Claims, 13 Drawing Sheets**



**Related U.S. Application Data**

(60) Provisional application No. 62/464,289, filed on Feb. 27, 2017.

(51) **Int. Cl.**

**G06F 16/27** (2019.01)  
**G06F 16/951** (2019.01)  
**G06F 16/22** (2019.01)  
**G06F 16/2458** (2019.01)

(52) **U.S. Cl.**

CPC ..... **G06F 16/2282** (2019.01); **G06F 16/2477** (2019.01); **G06F 16/27** (2019.01); **G06F 16/278** (2019.01); **G06F 16/951** (2019.01)

(56)

**References Cited**

U.S. PATENT DOCUMENTS

2008/0091691 A1 4/2008 Tsuji  
 2008/0162409 A1 7/2008 Meijer et al.  
 2010/0318587 A1 12/2010 Seet et al.

2012/0254174 A1 10/2012 Mitra et al.  
 2012/0254175 A1 10/2012 Horowitz et al.  
 2013/0080393 A1 3/2013 Bird et al.  
 2013/0103657 A1 4/2013 Ikawa et al.  
 2014/0172866 A1 6/2014 Lin et al.  
 2015/0134796 A1 5/2015 Theimer et al.  
 2015/0149413 A1 5/2015 Lee et al.  
 2015/0347555 A1\* 12/2015 Vyas ..... G06F 16/278  
 707/662  
 2016/0085839 A1\* 3/2016 D'Halluin ..... G06F 16/278  
 707/747  
 2016/0350392 A1 12/2016 Rice et al.  
 2017/0046412 A1 2/2017 Levy et al.  
 2017/0322996 A1 11/2017 Horowitz et al.  
 2017/0344441 A1\* 11/2017 Horowitz ..... G06F 11/1474  
 2017/0371910 A1 12/2017 Joo et al.  
 2019/0121902 A1\* 4/2019 Ryan ..... G06F 16/278

OTHER PUBLICATIONS

PCT Invitation to Pay Additional Fees, PCT Application No. PCT/US2018/019990, Apr. 3, 2018, 2 pages.

\* cited by examiner

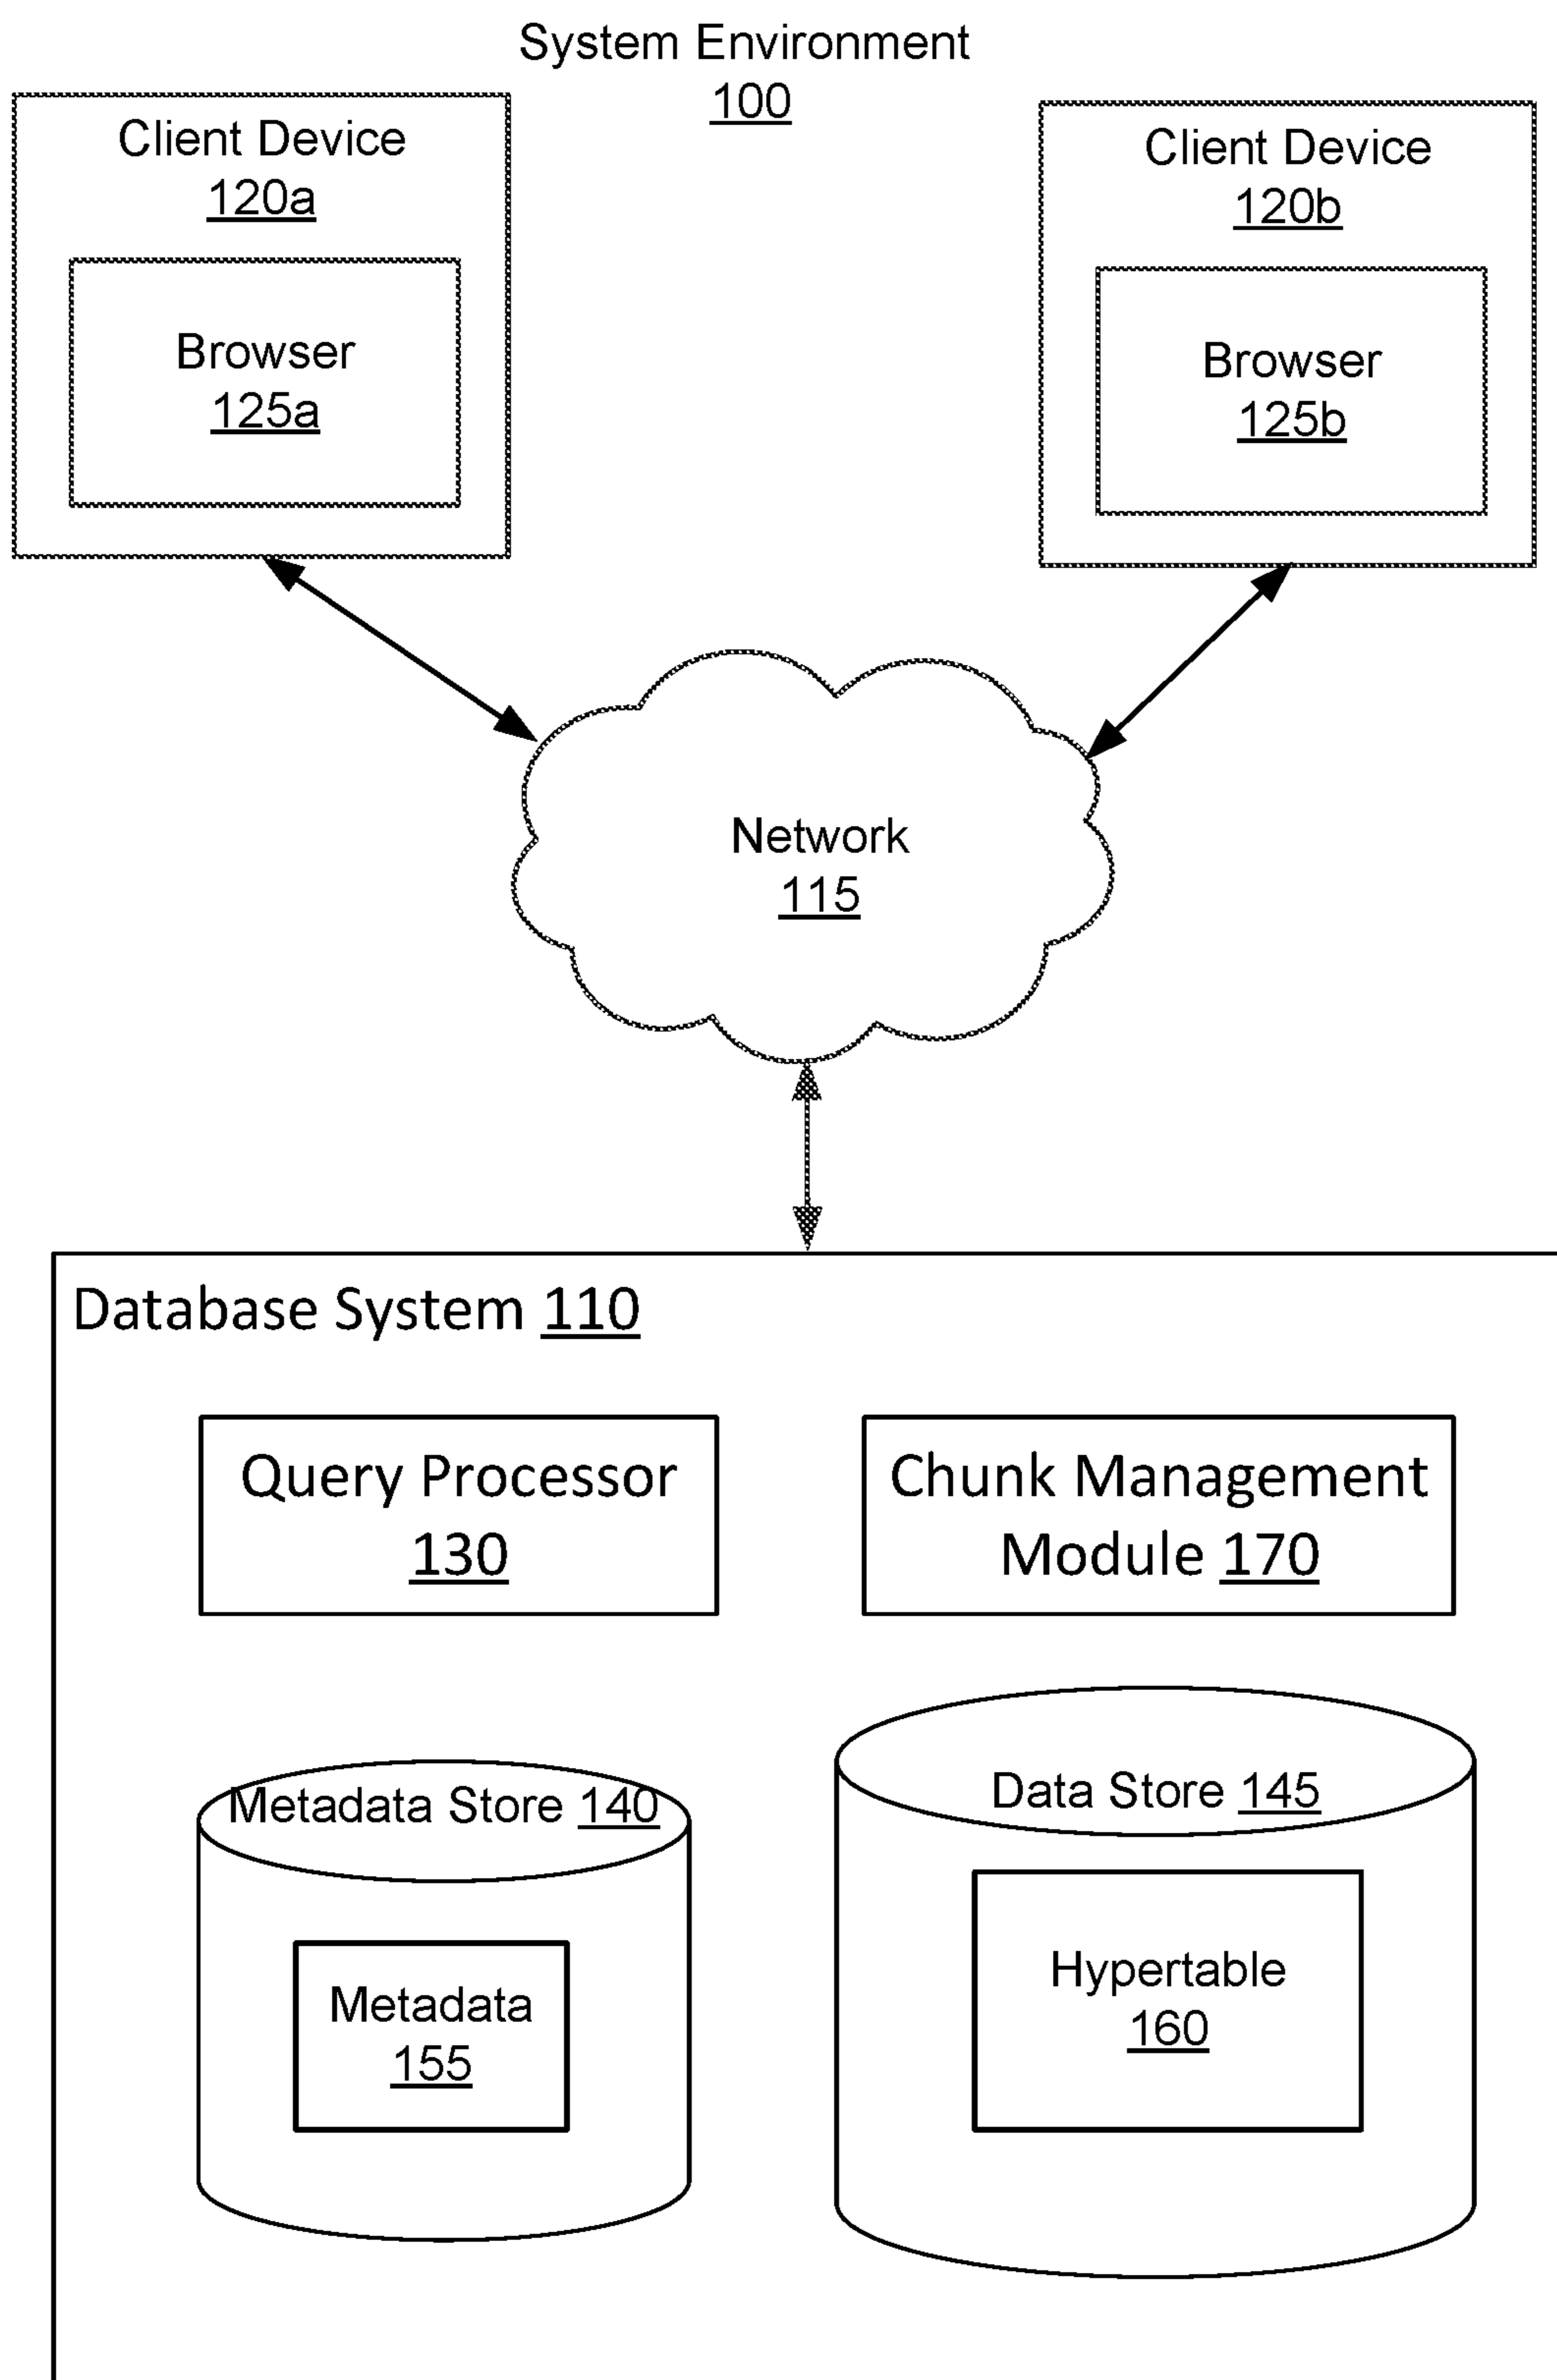


FIG. 1




FIG. 2




FIG. 3

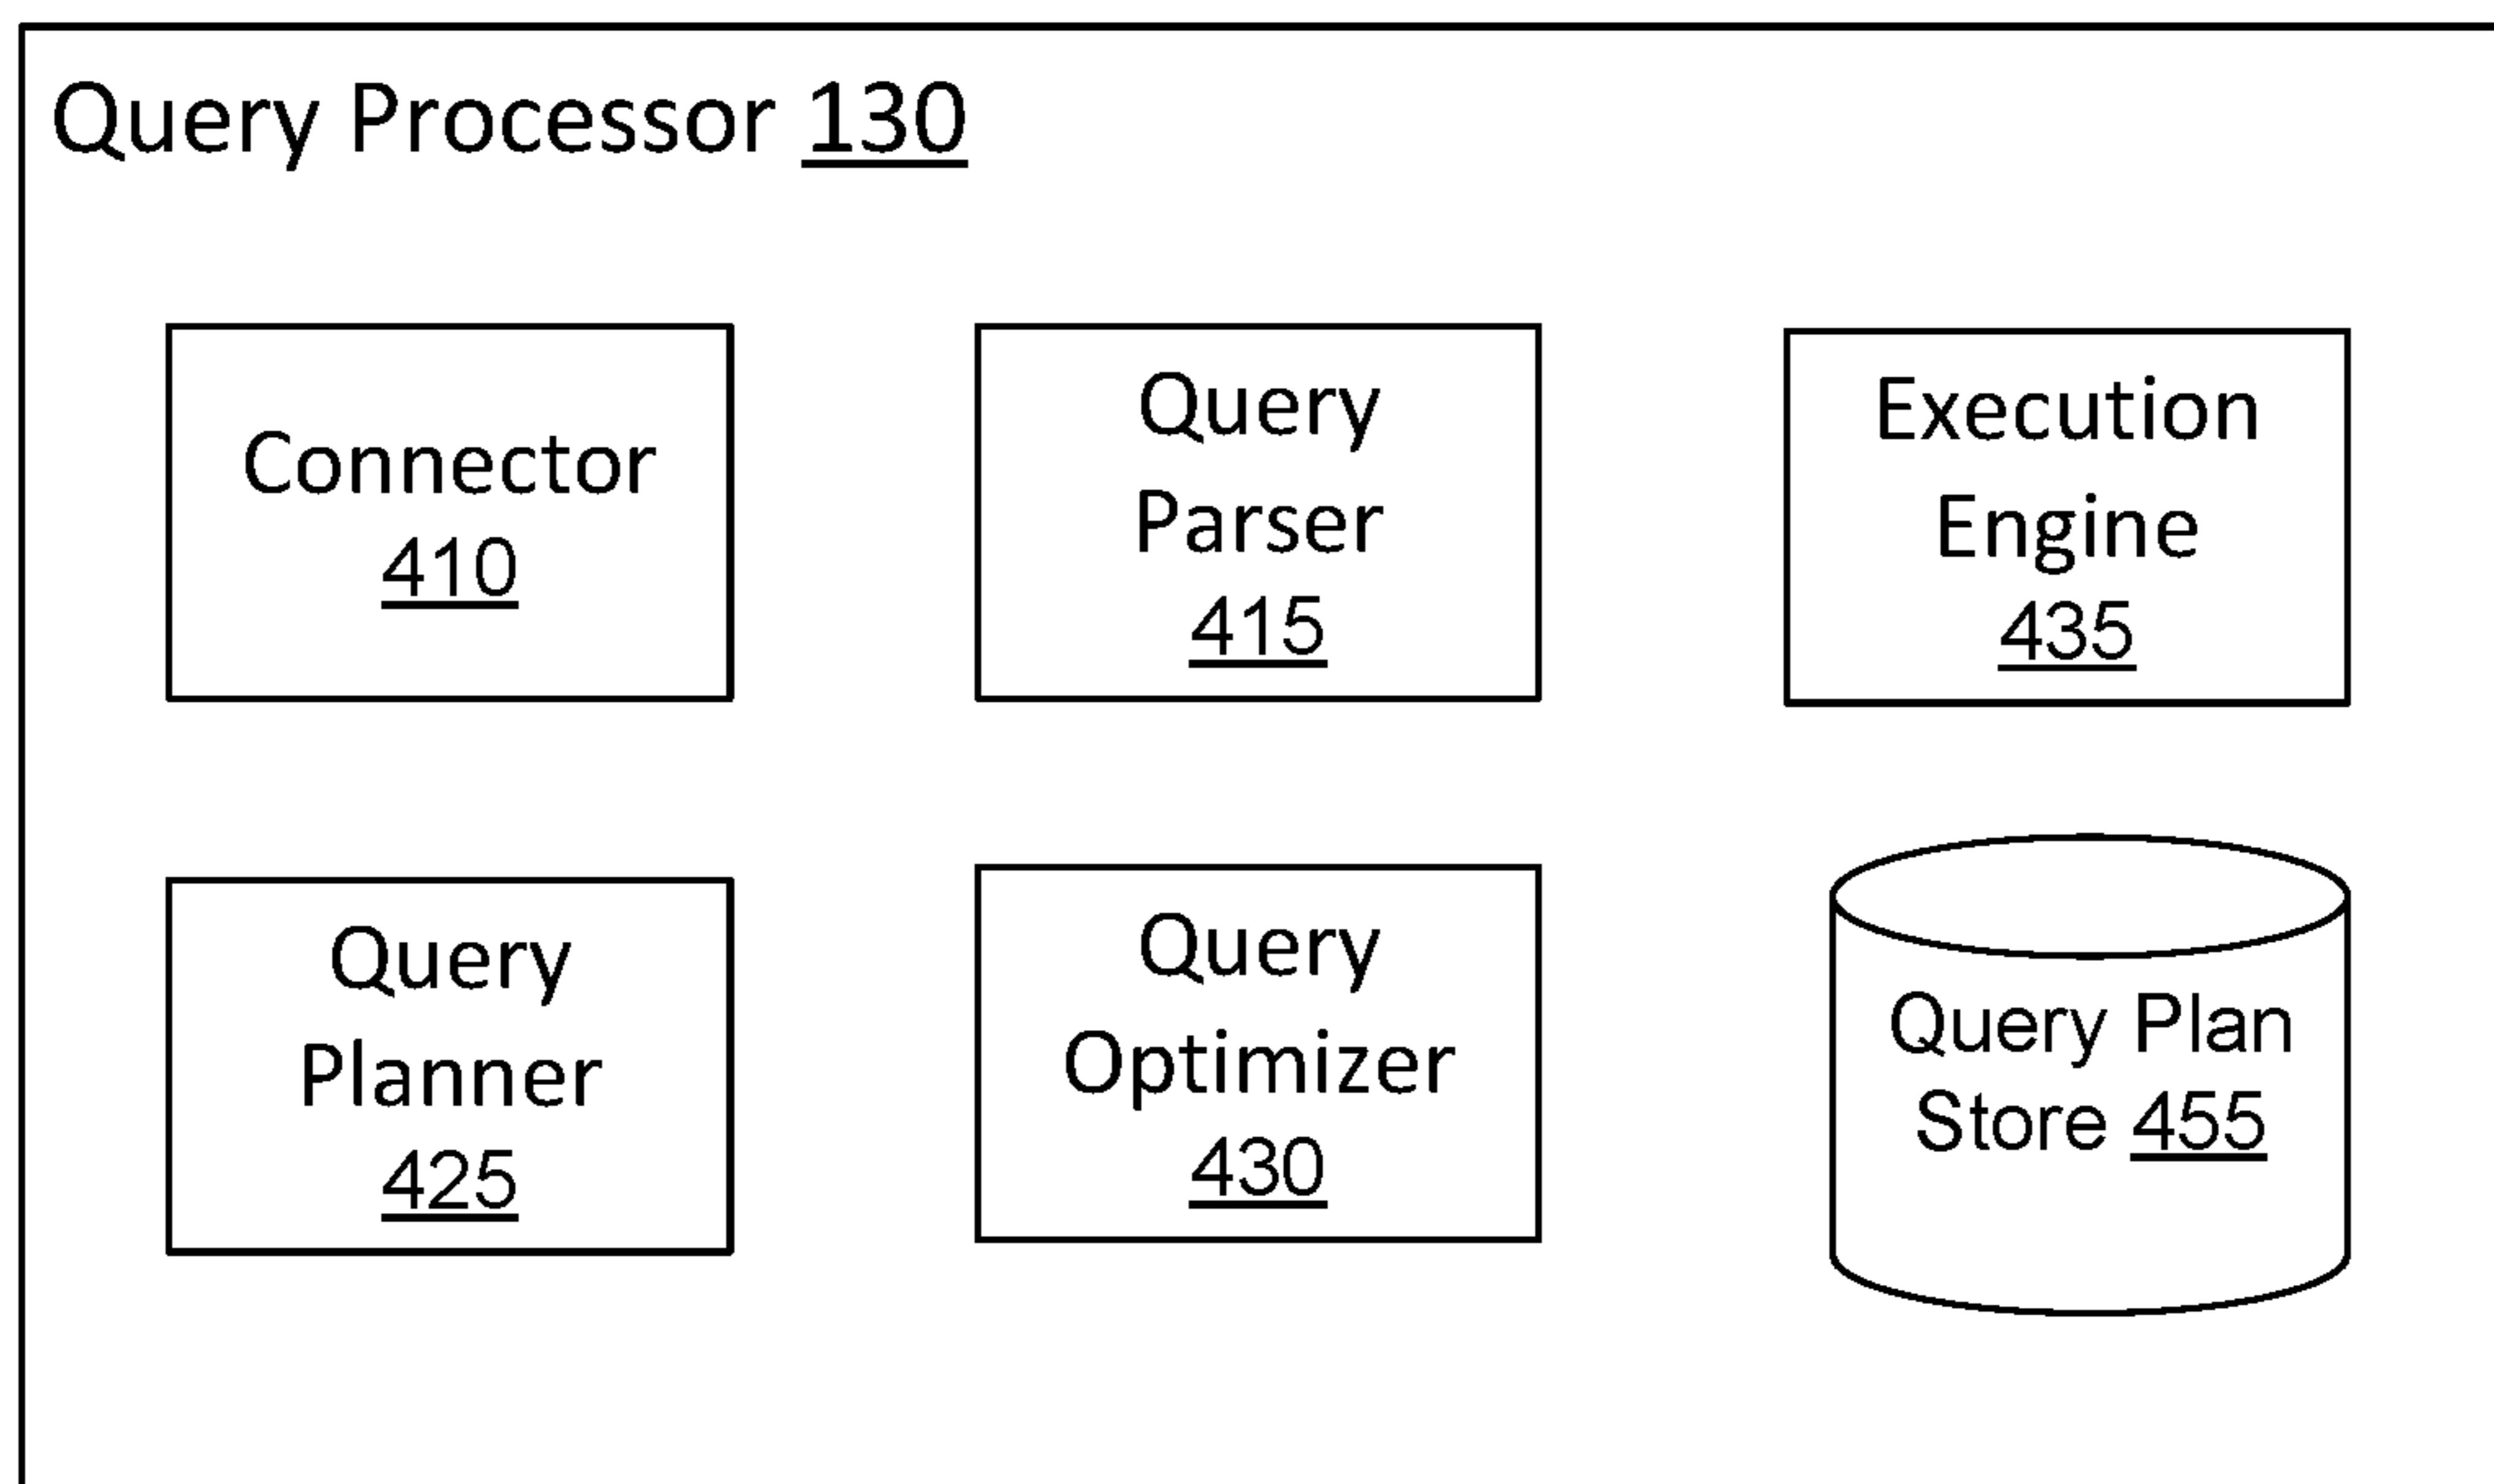


FIG. 4A

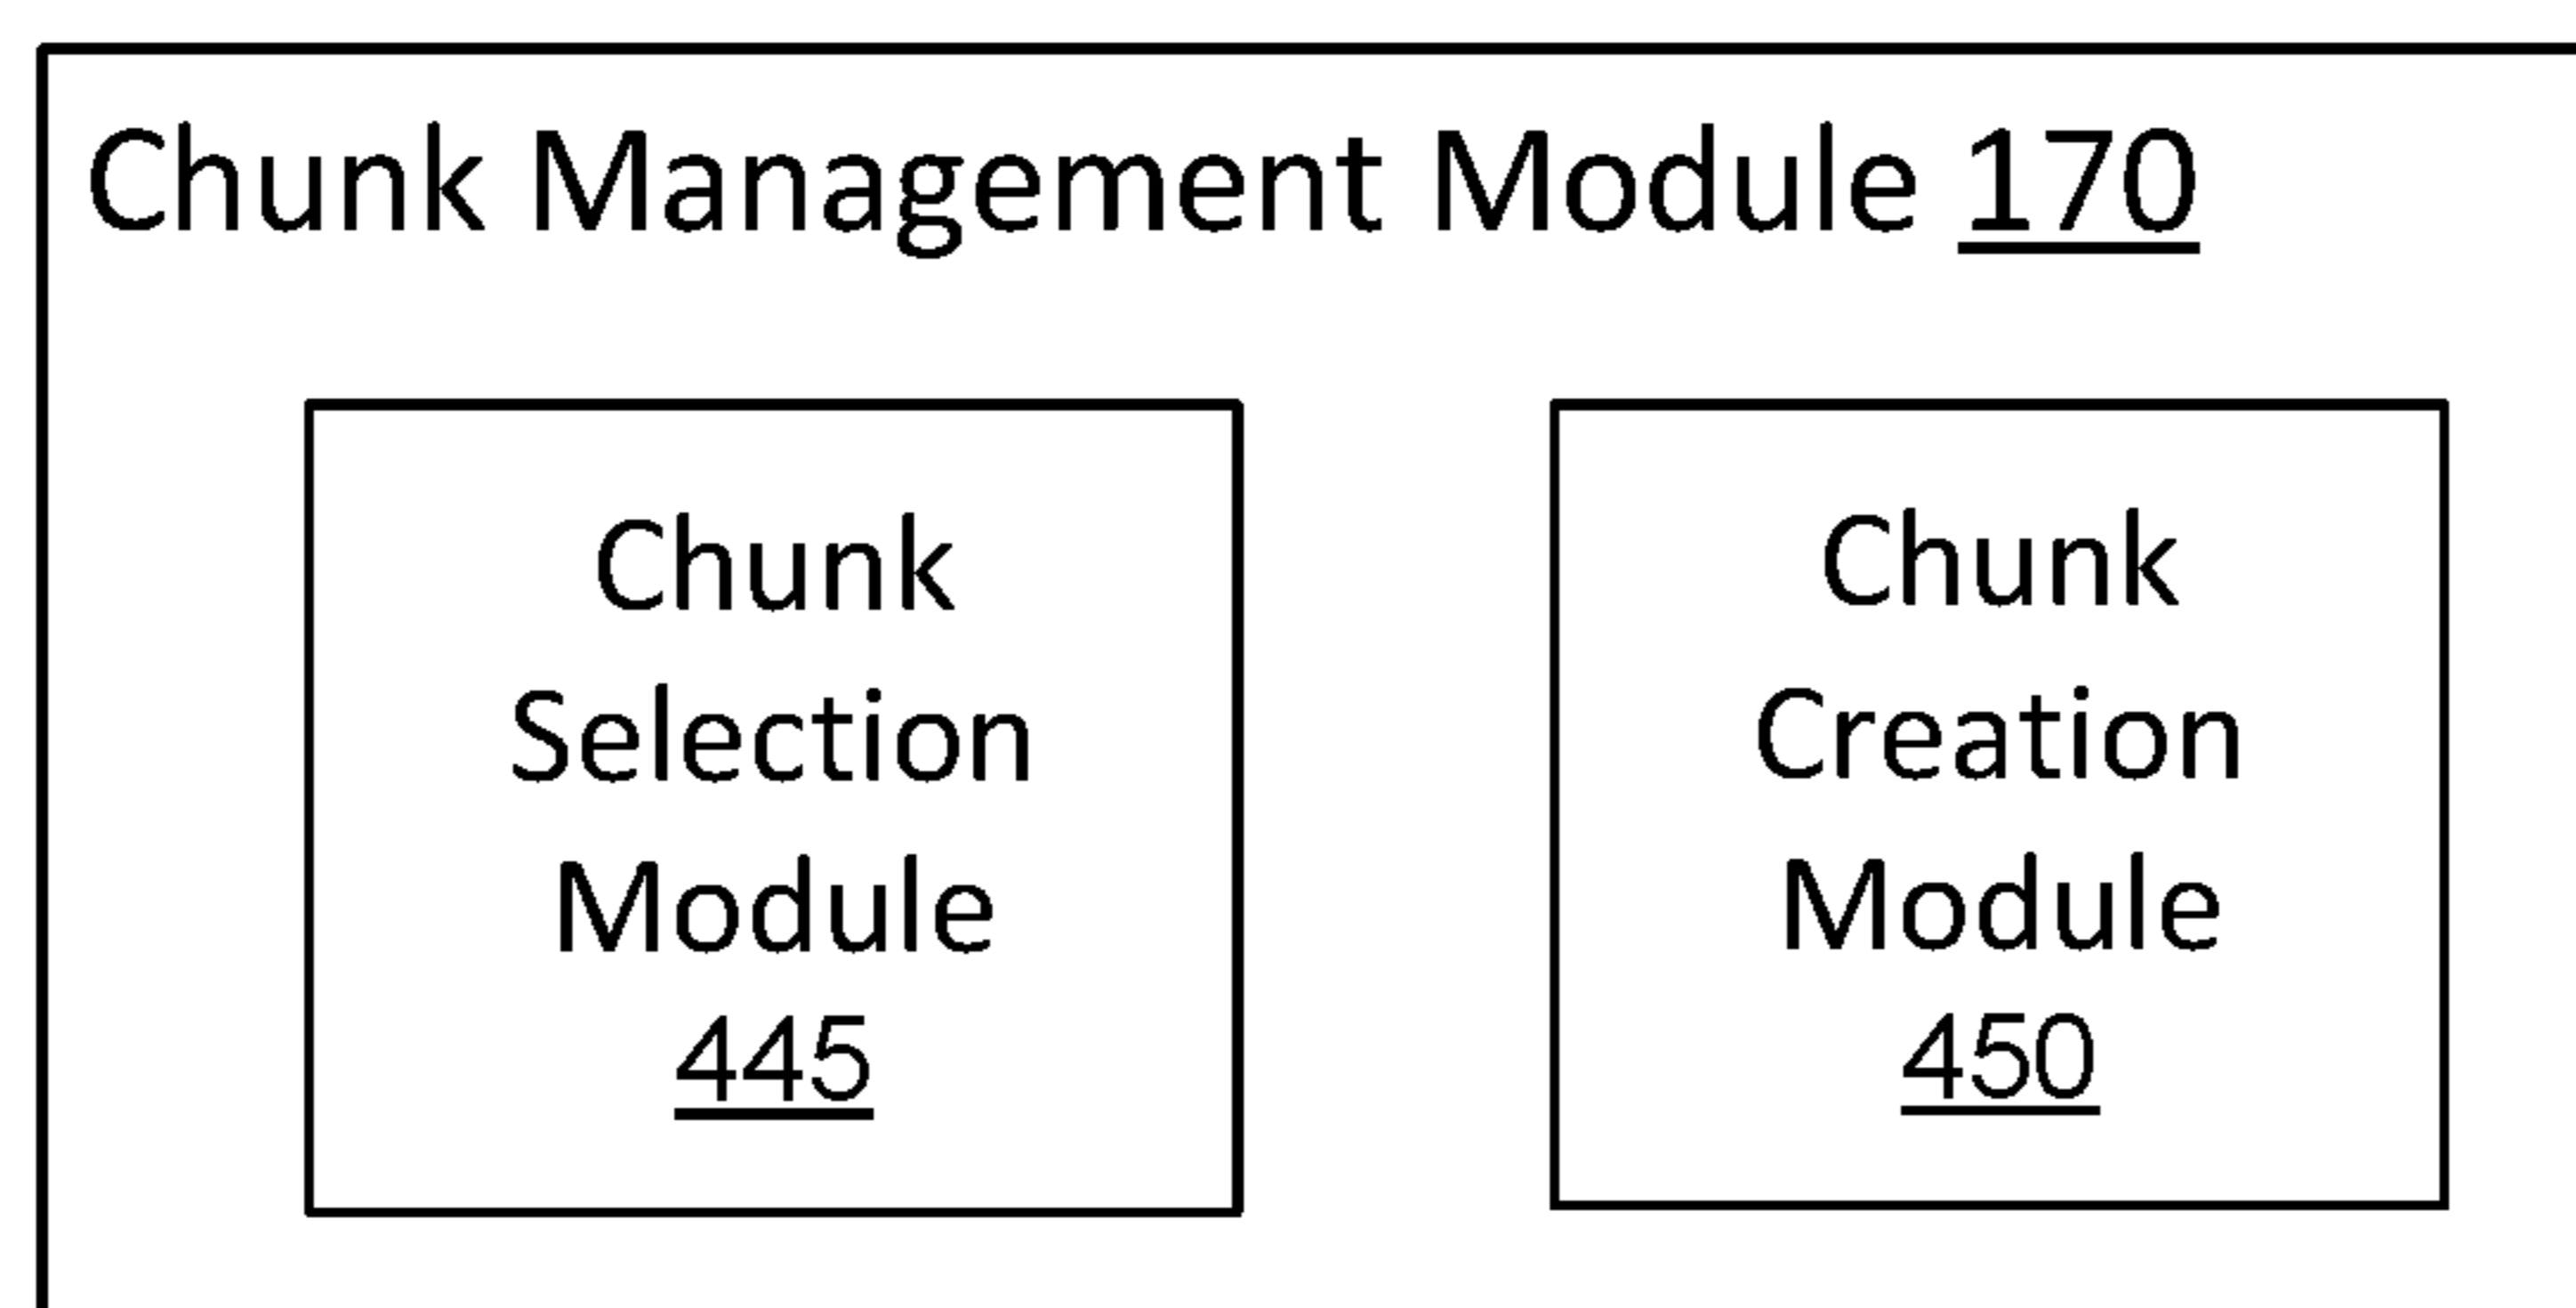


FIG. 4B

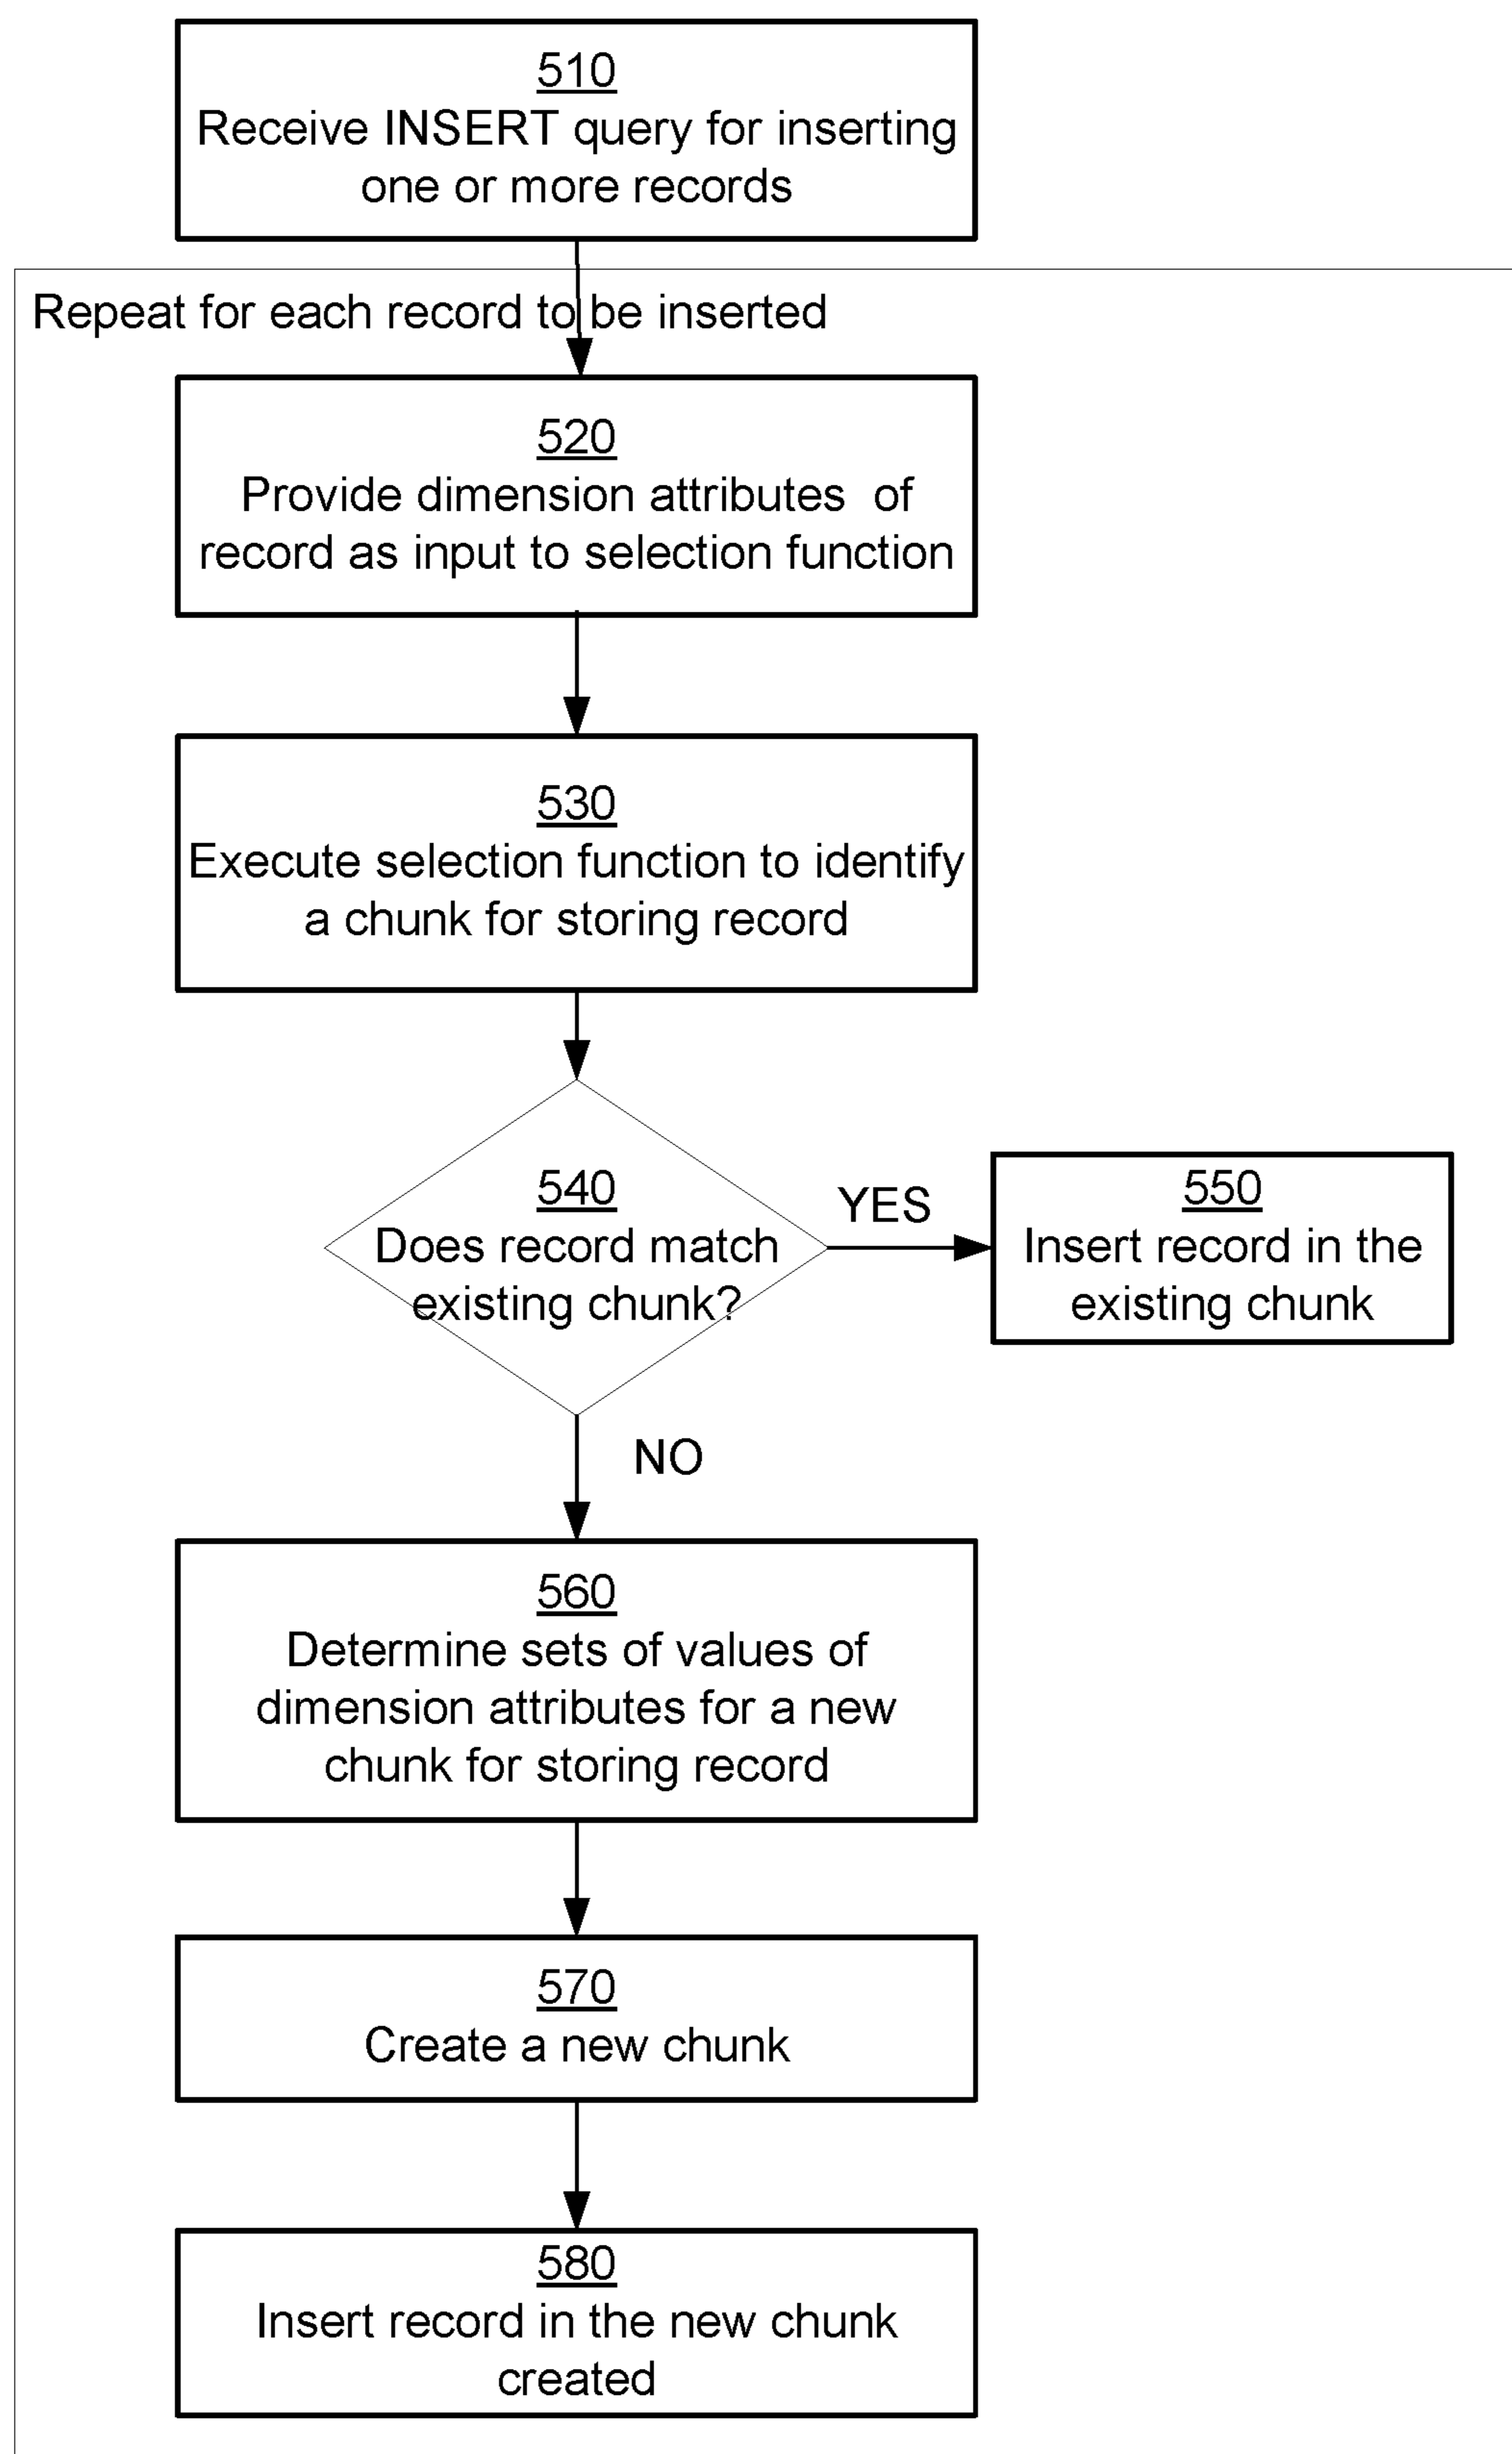


FIG. 5

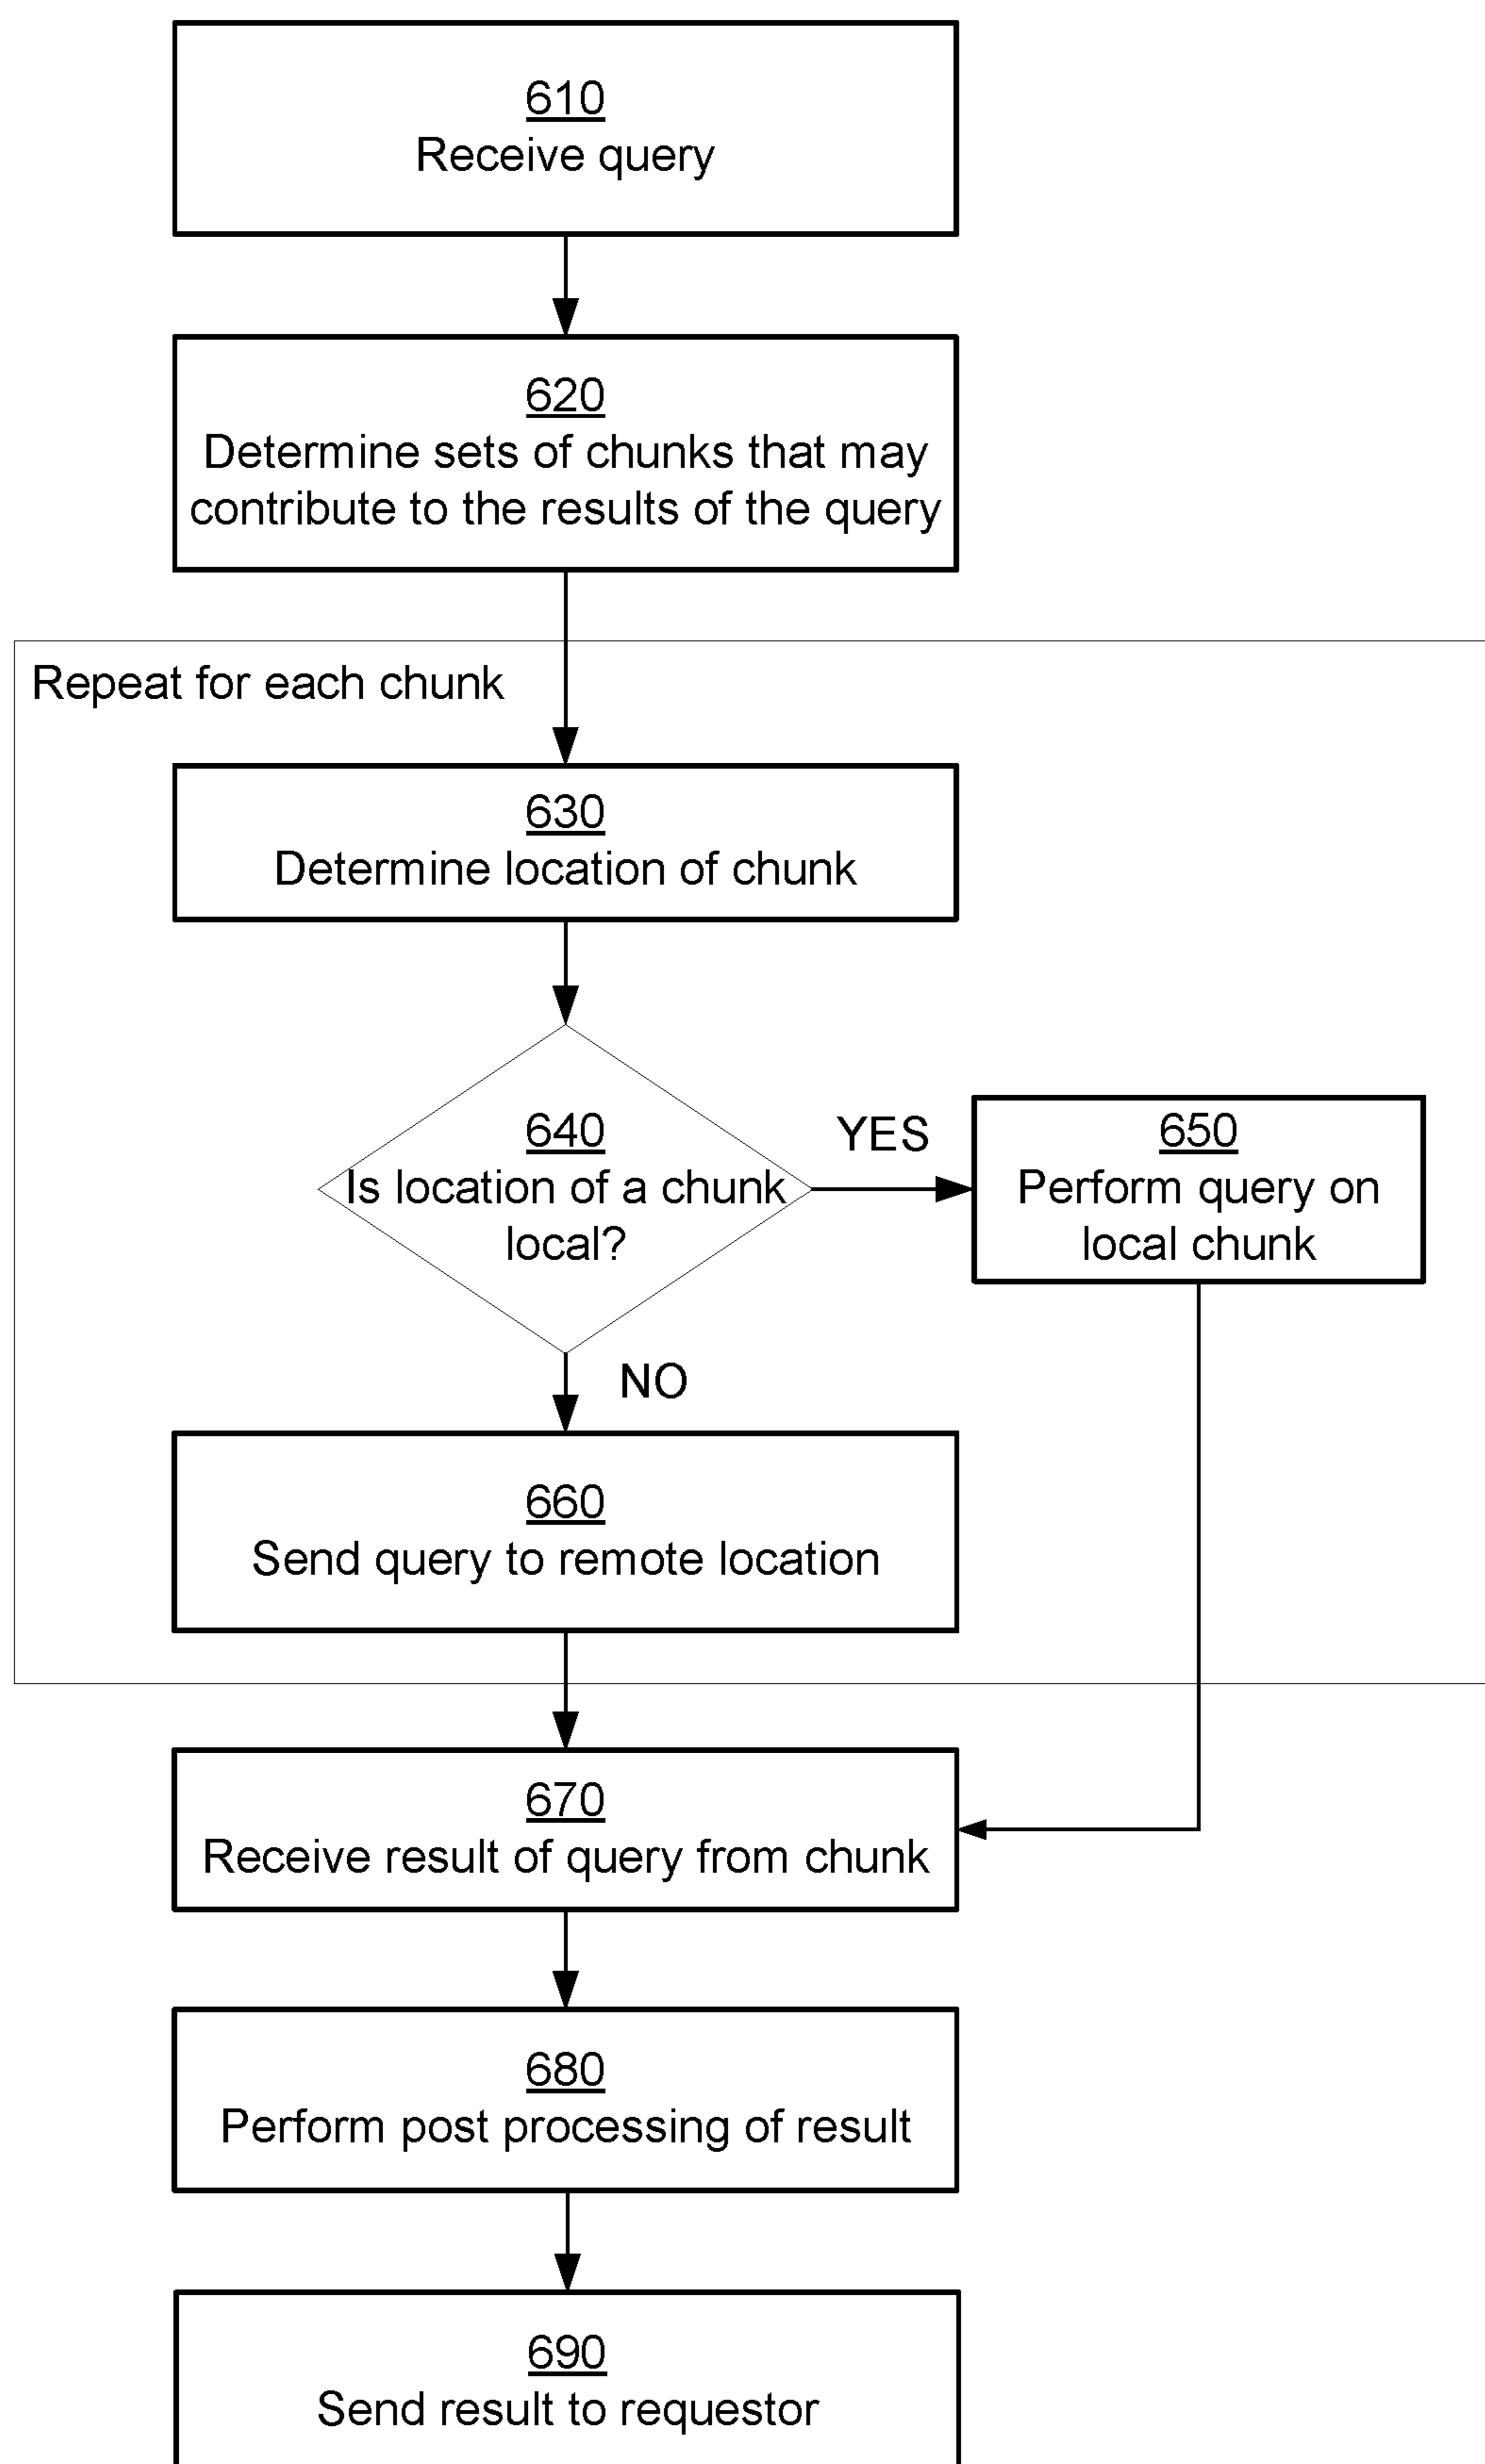


FIG. 6






FIG. 7(A)




FIG. 7(B)

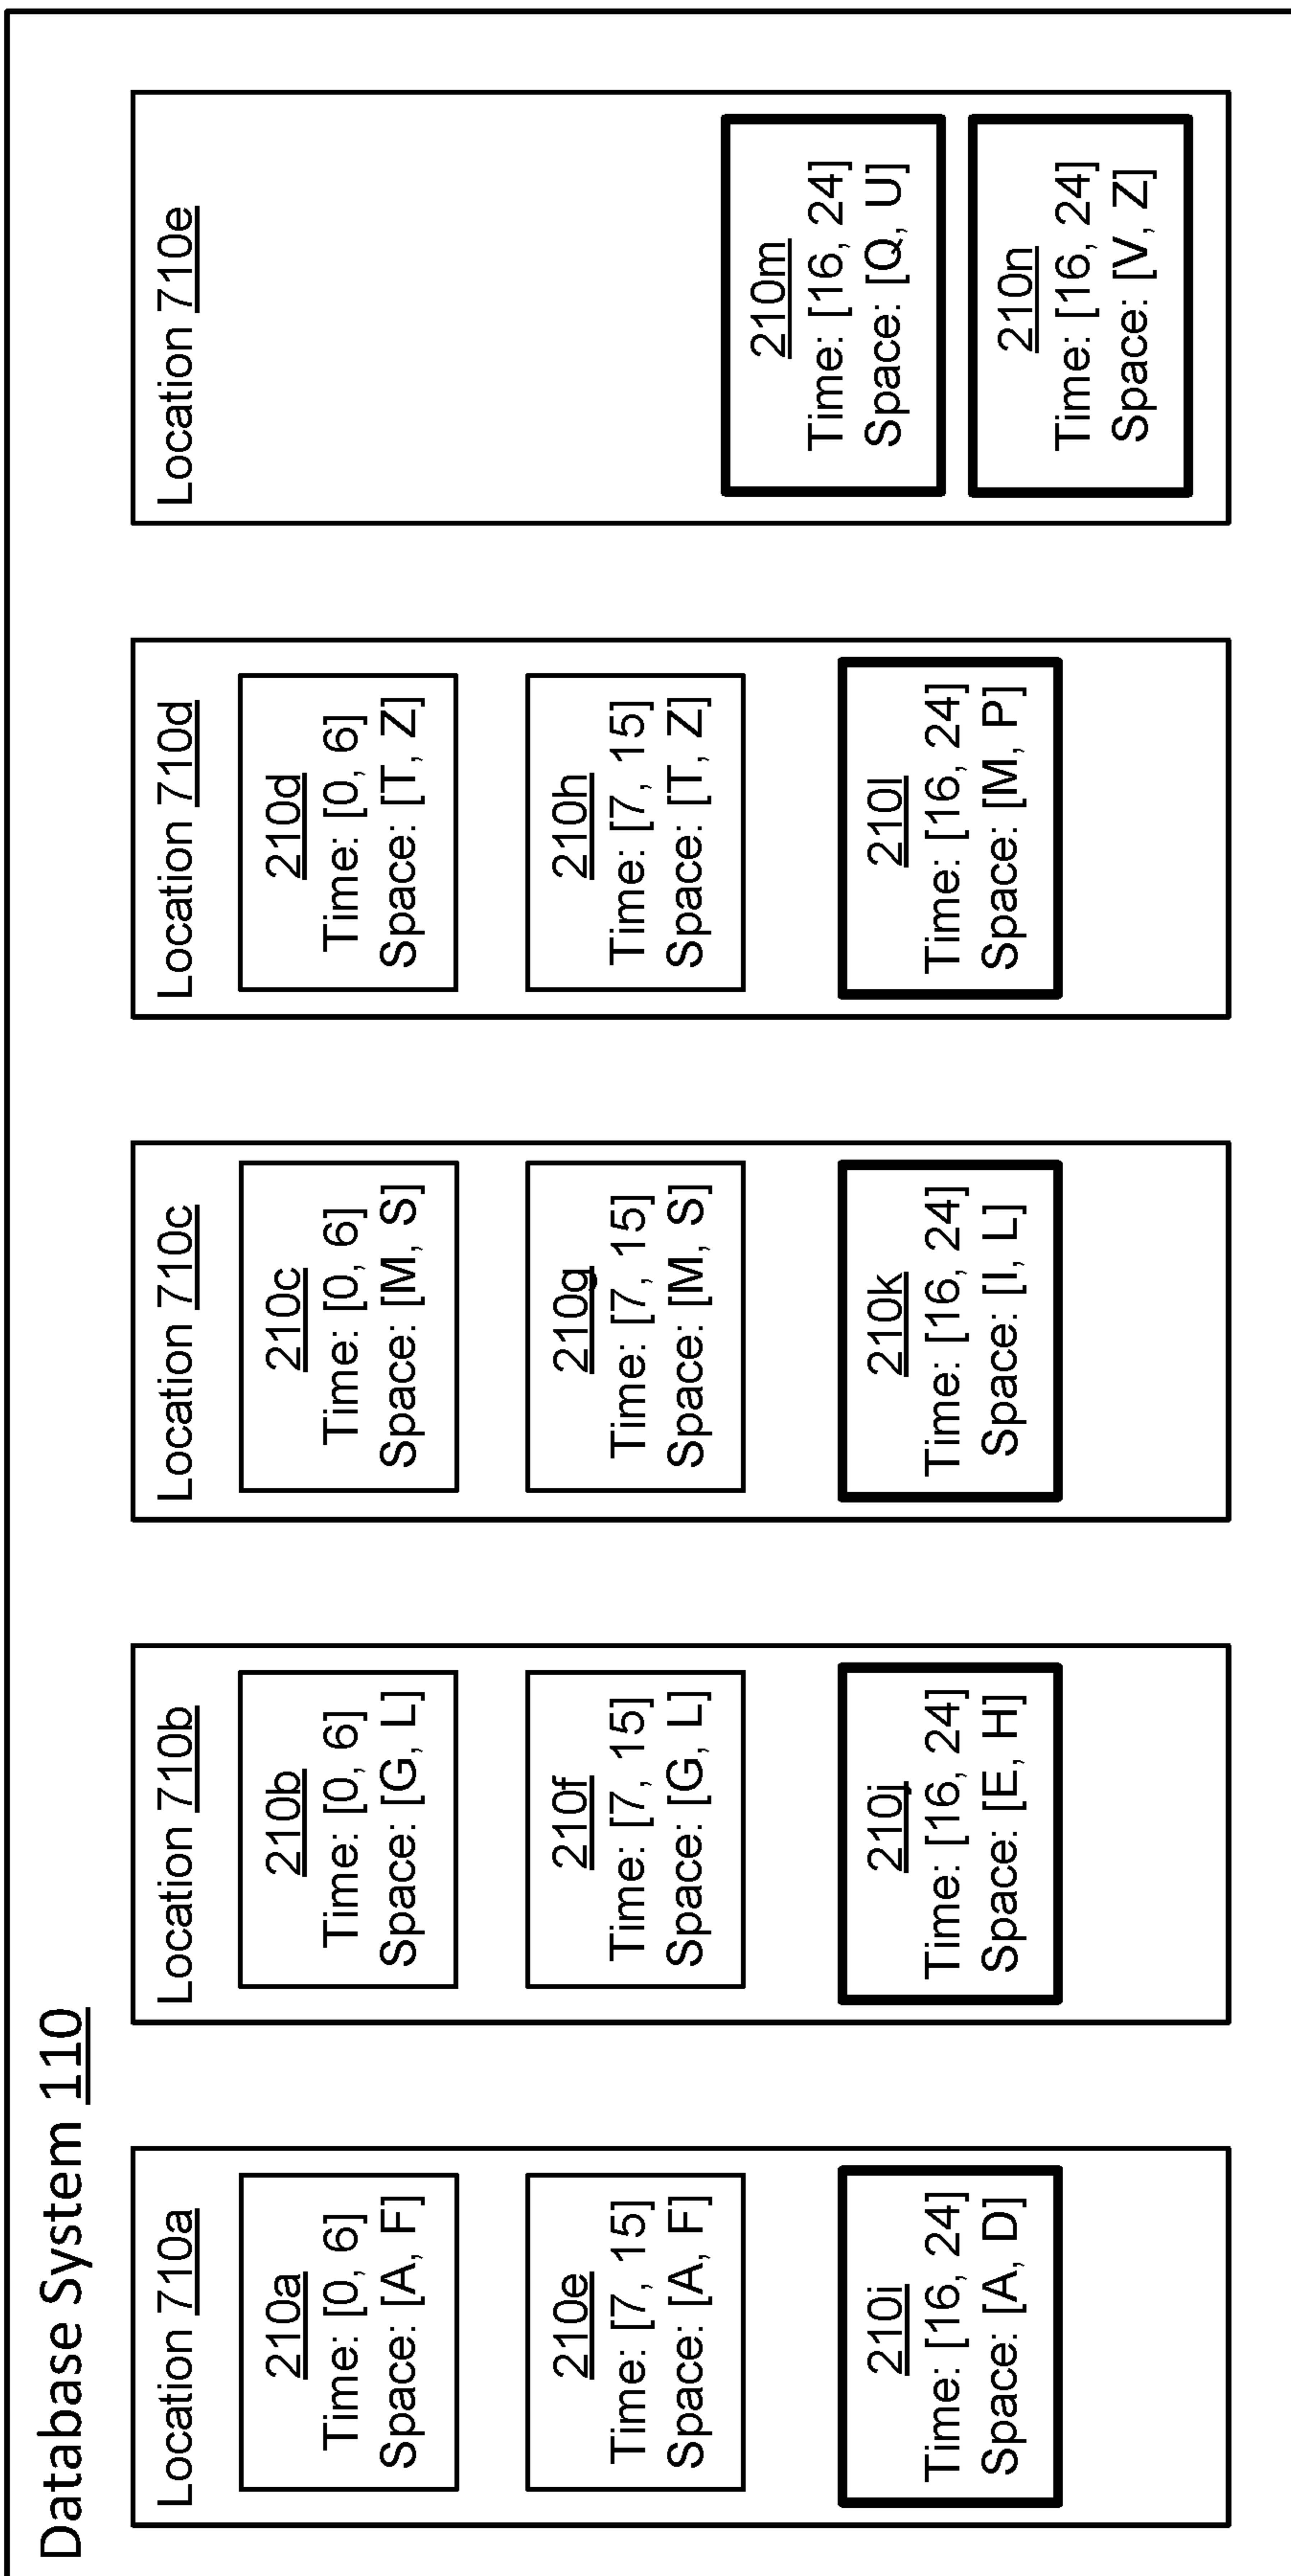


FIG. 7(C)

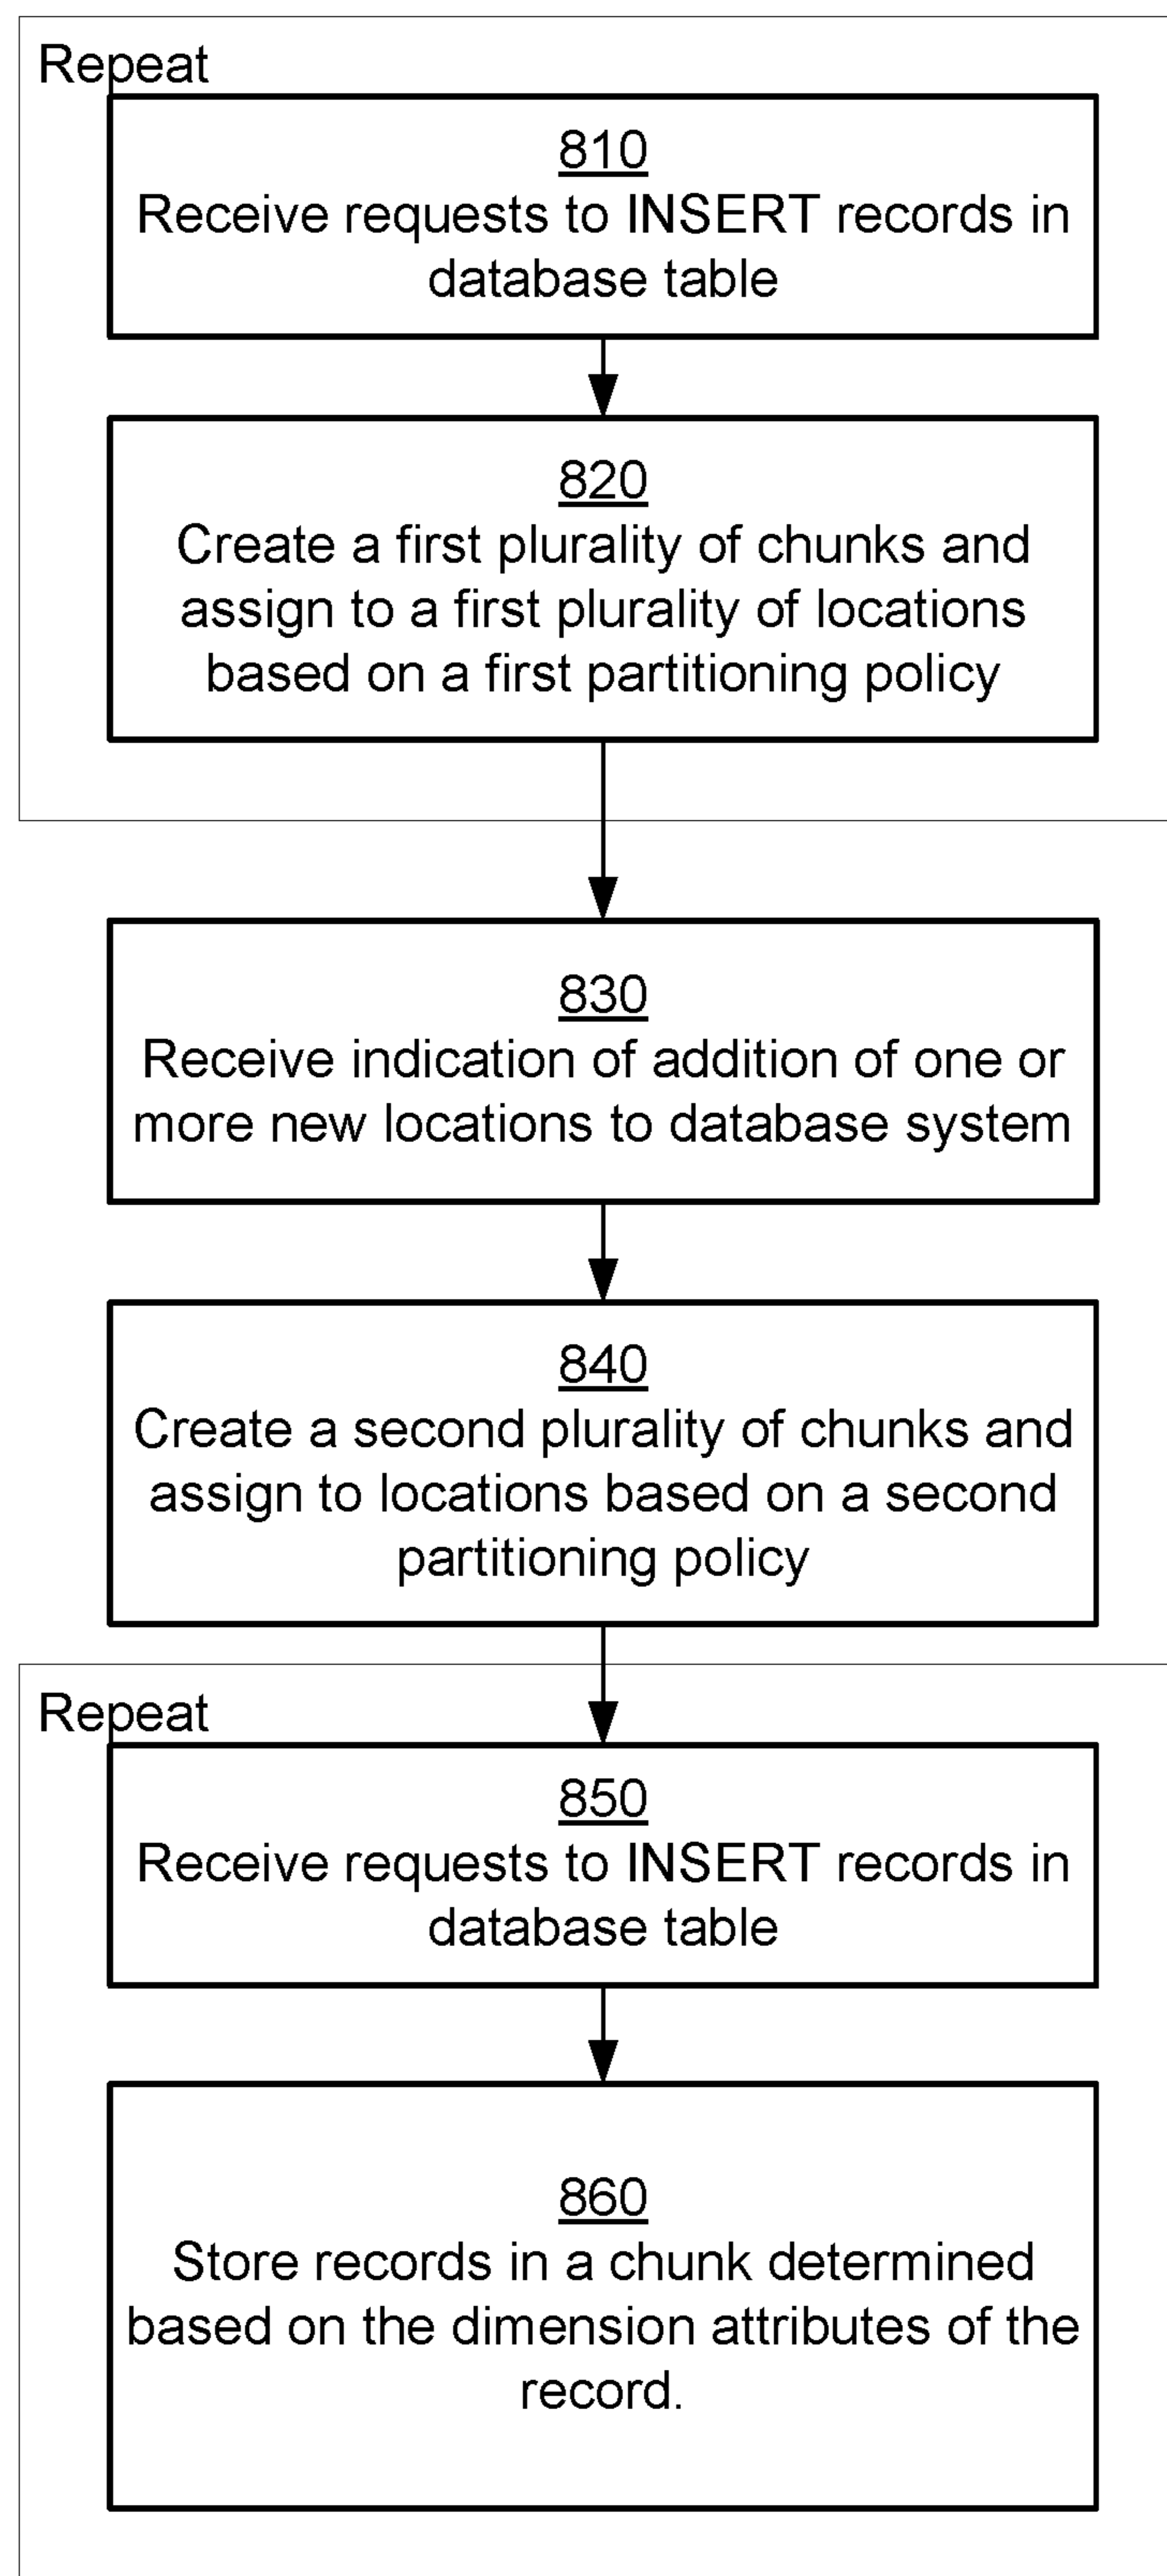


FIG. 8

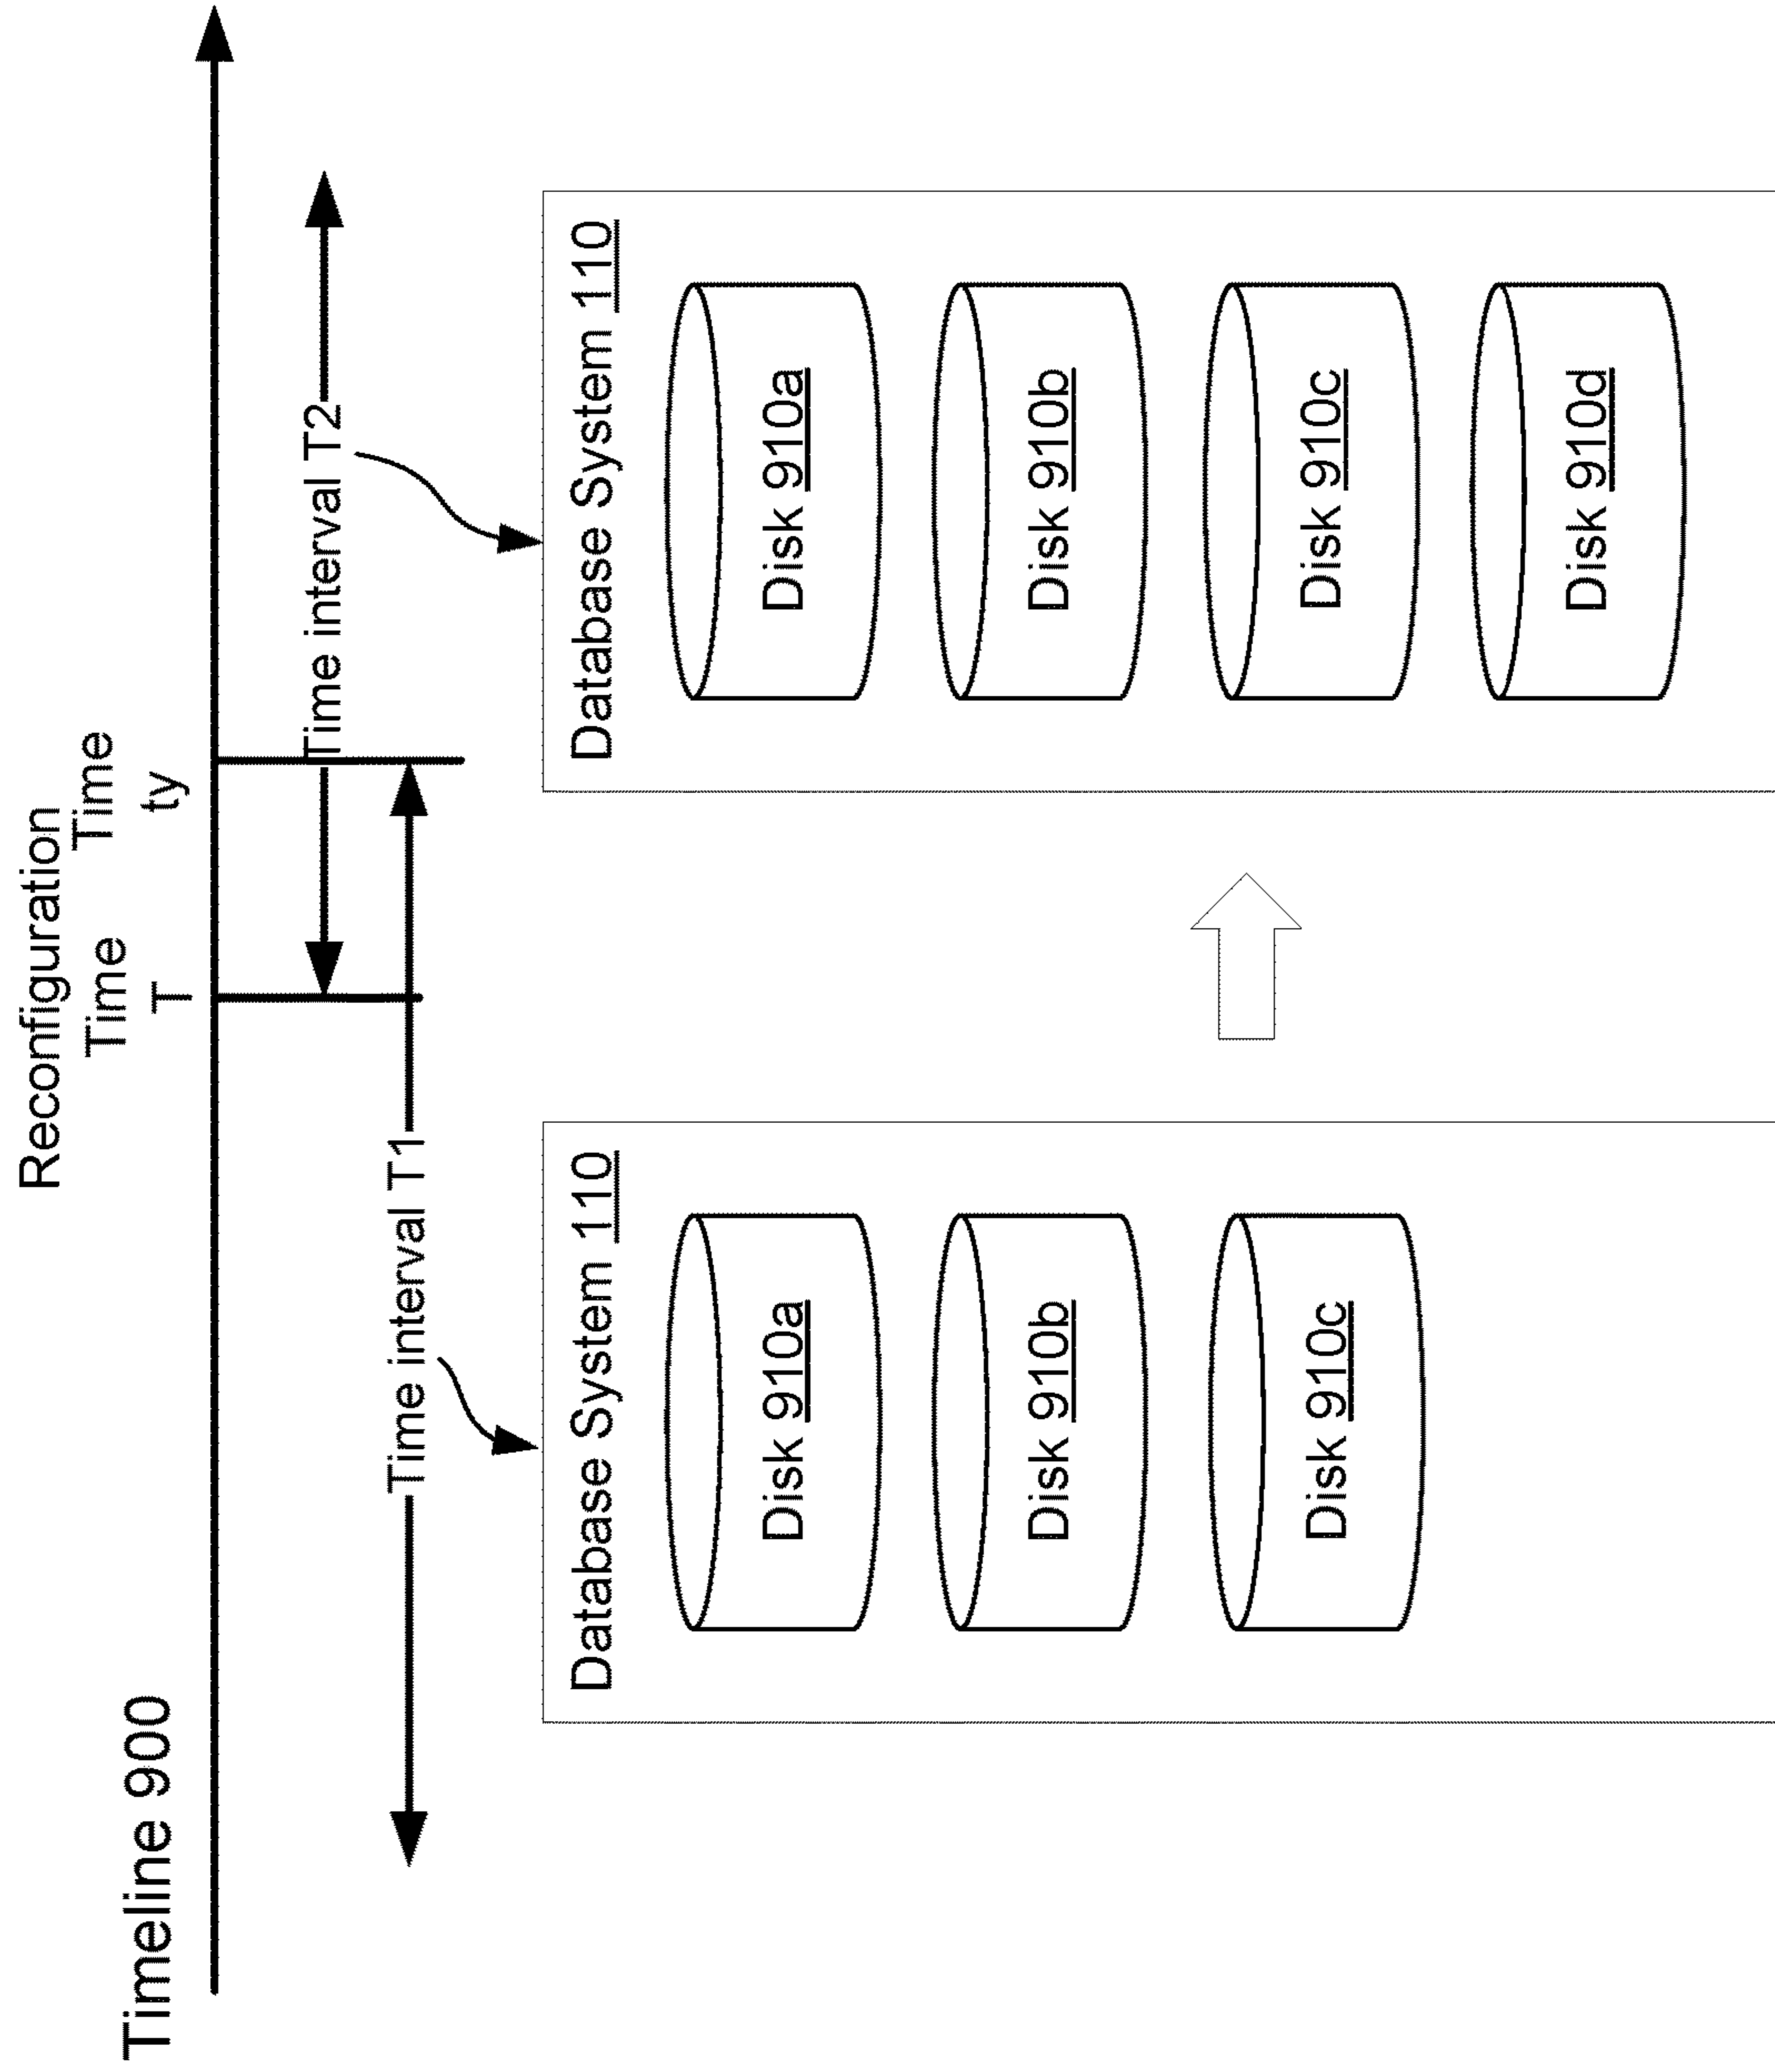


FIG. 9

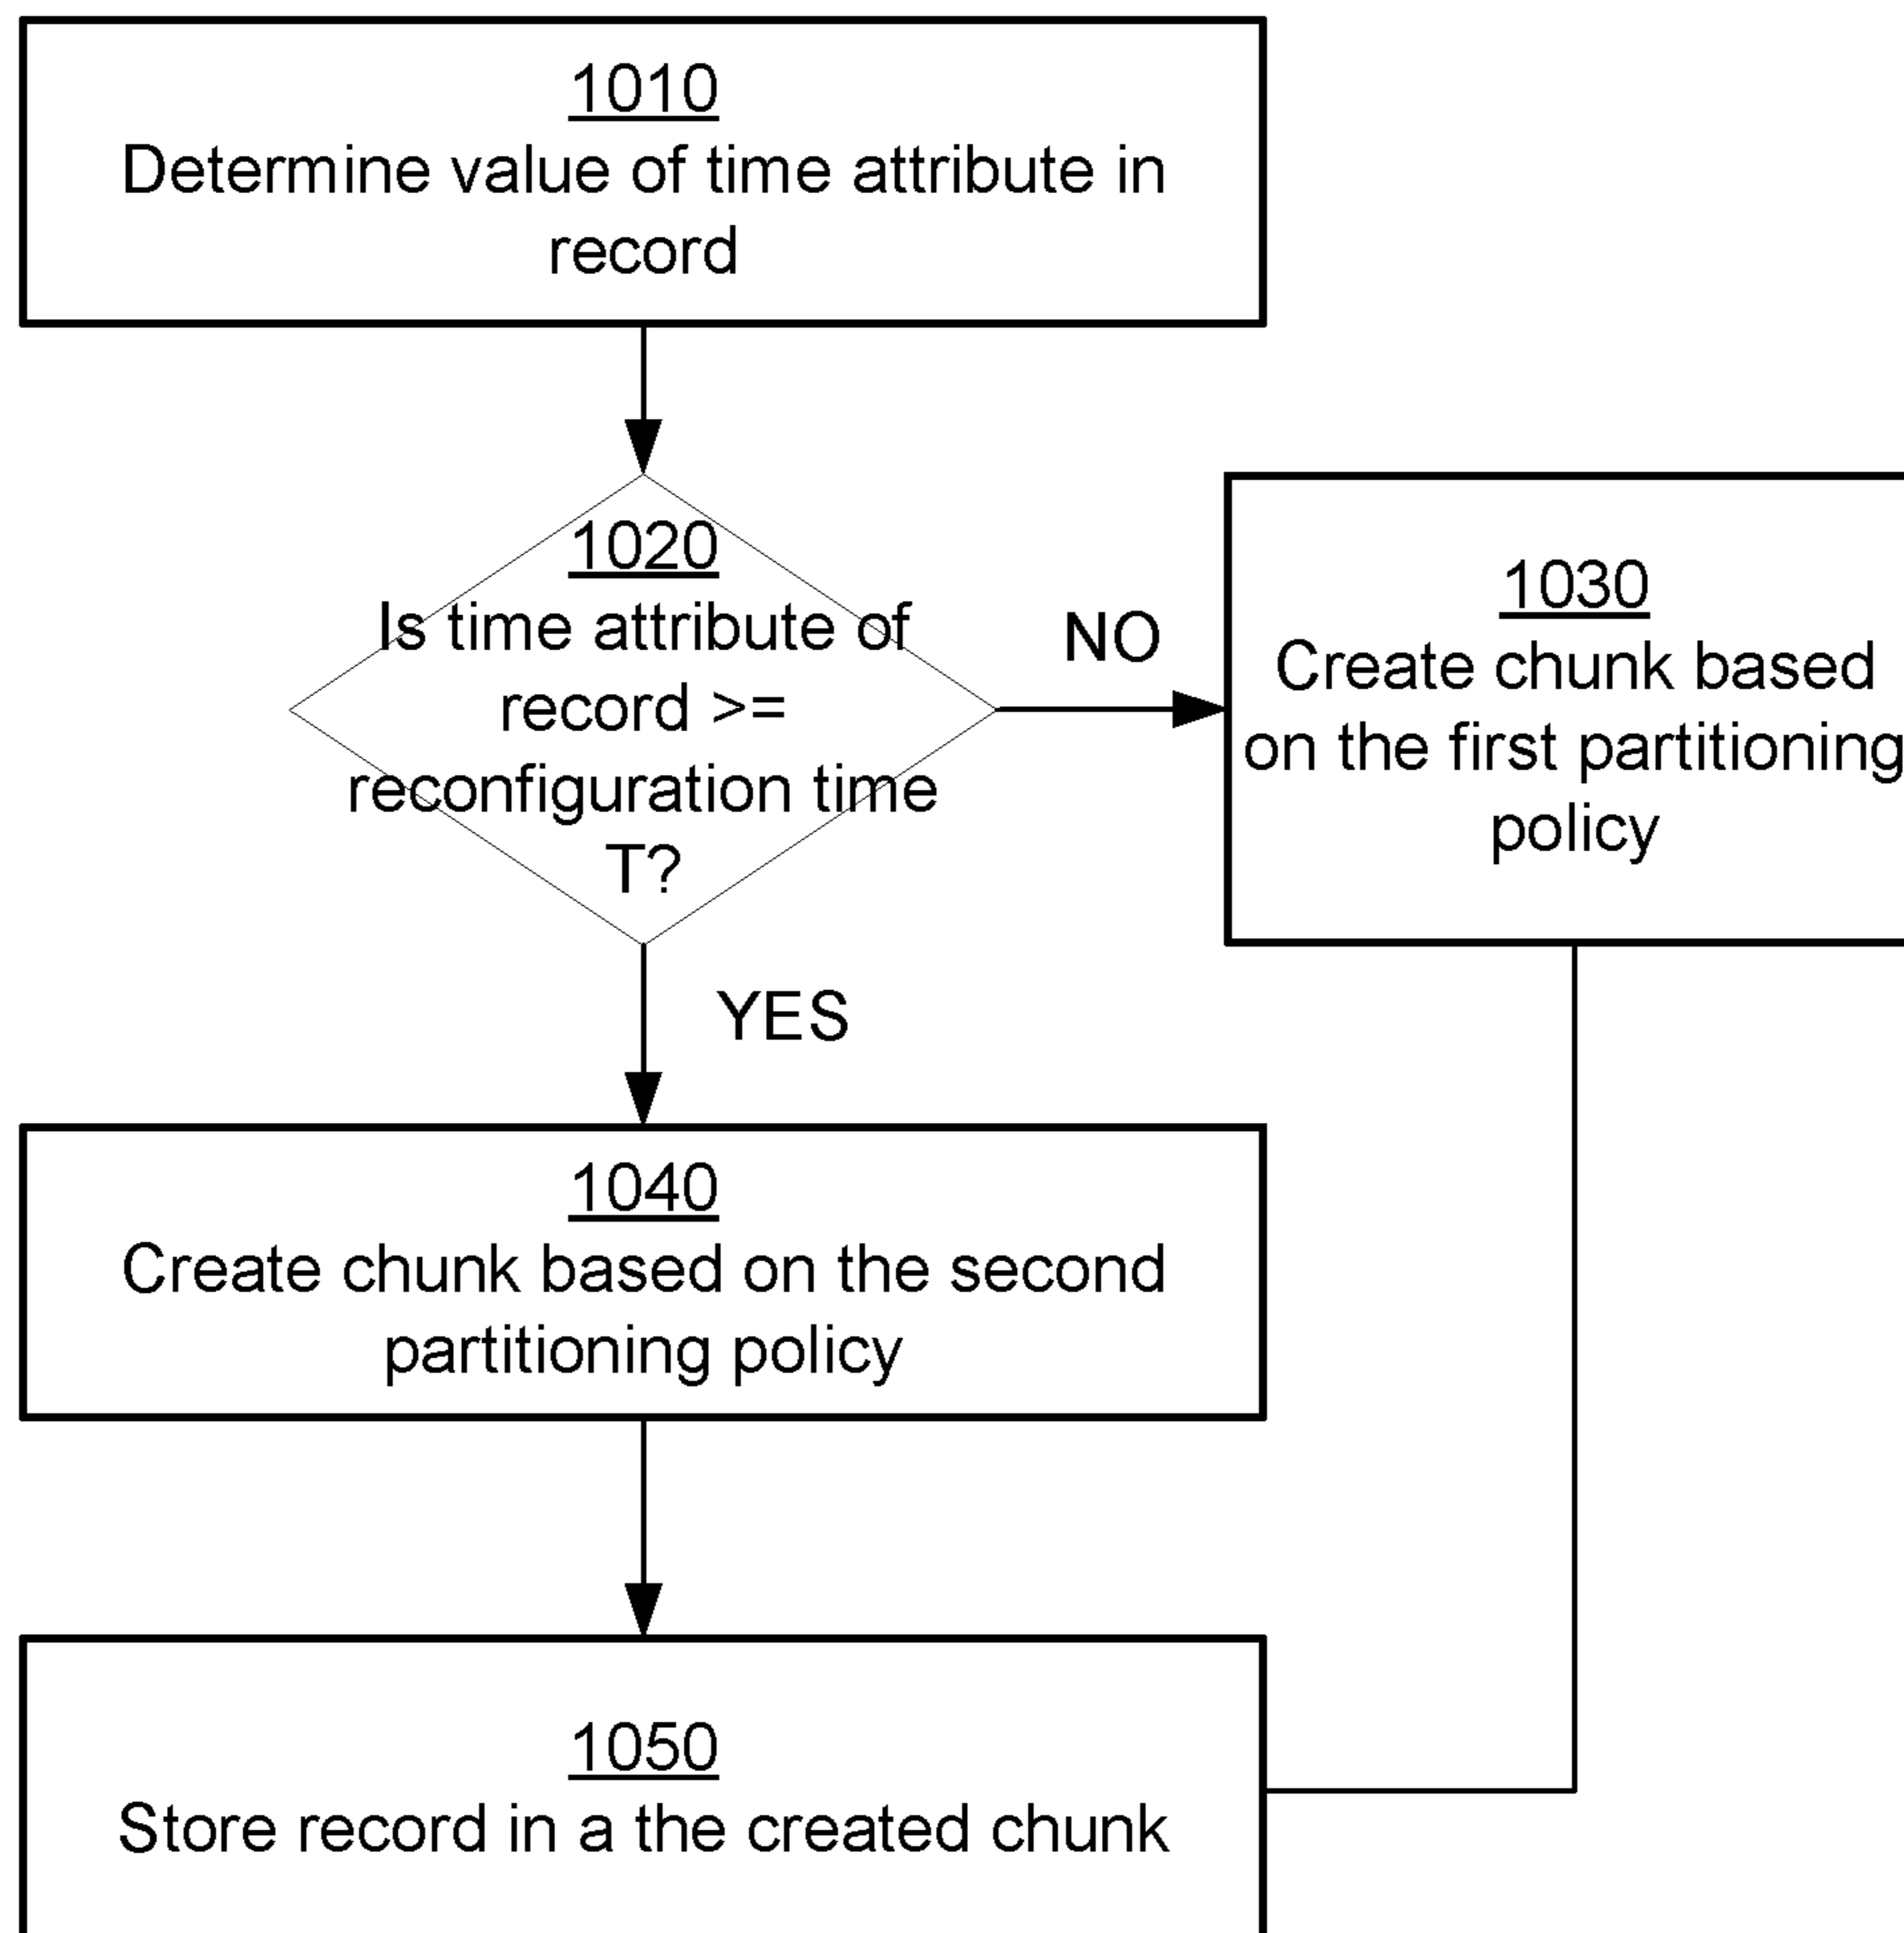


FIG. 10

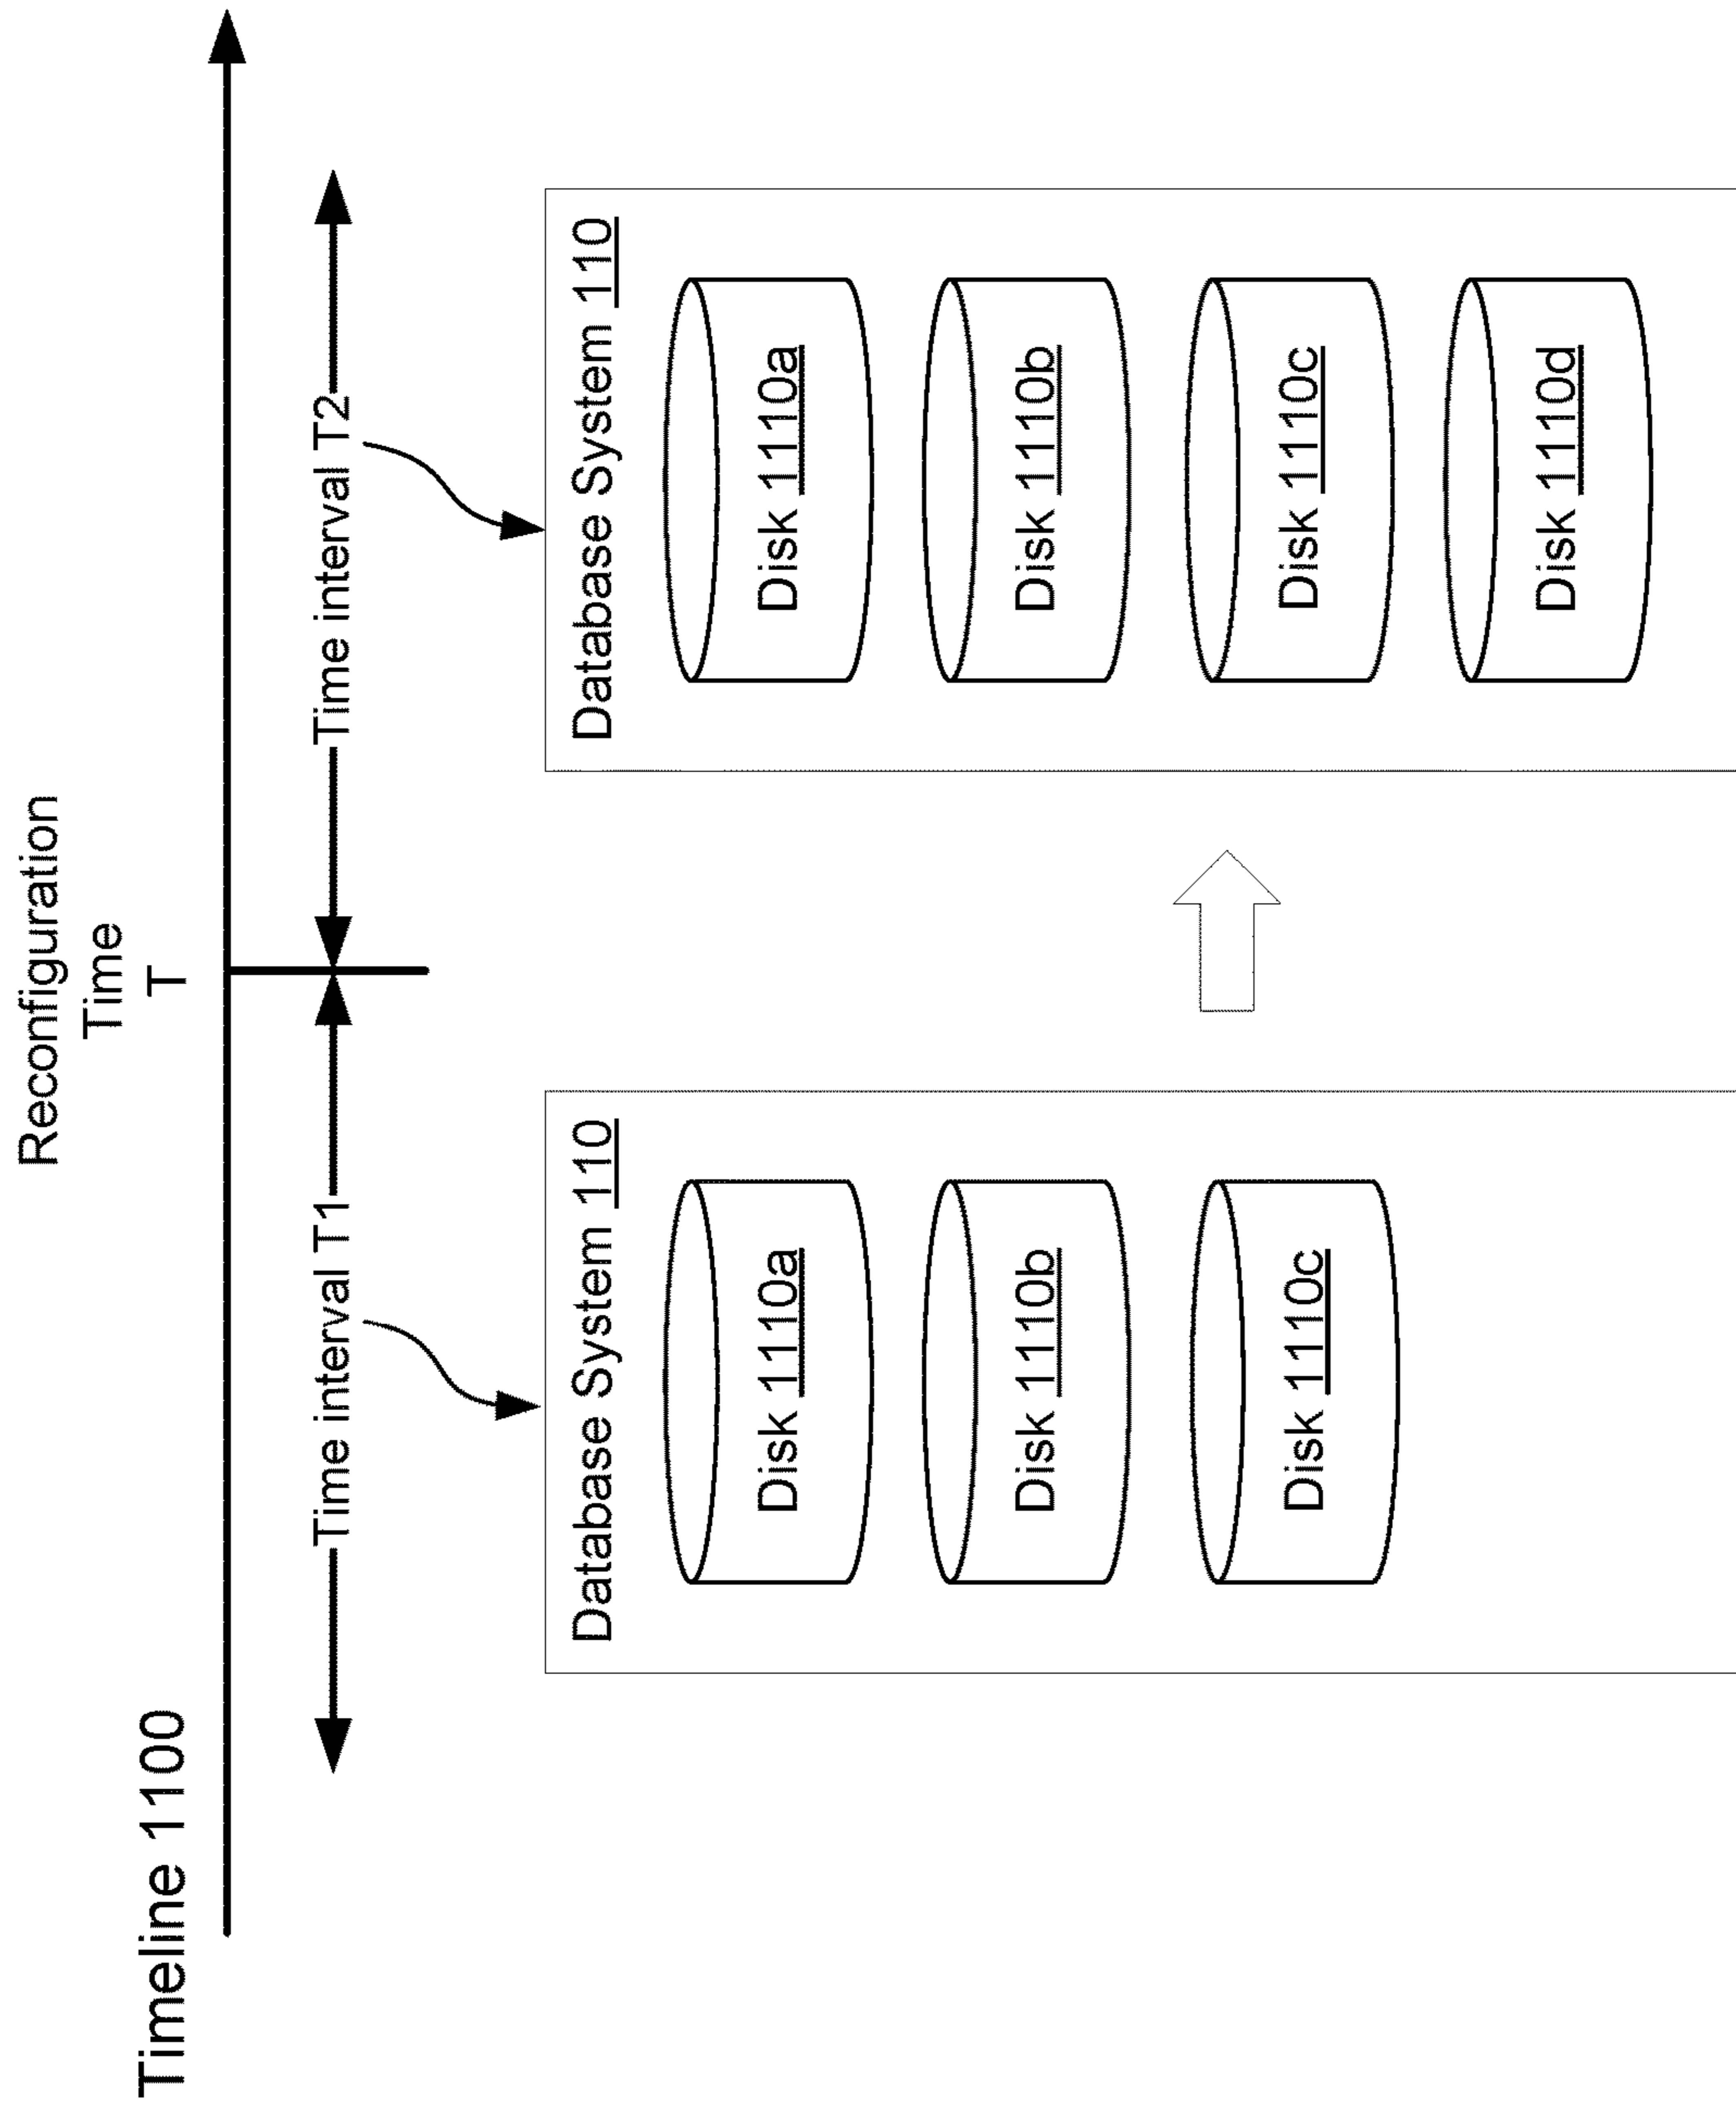


FIG. 11

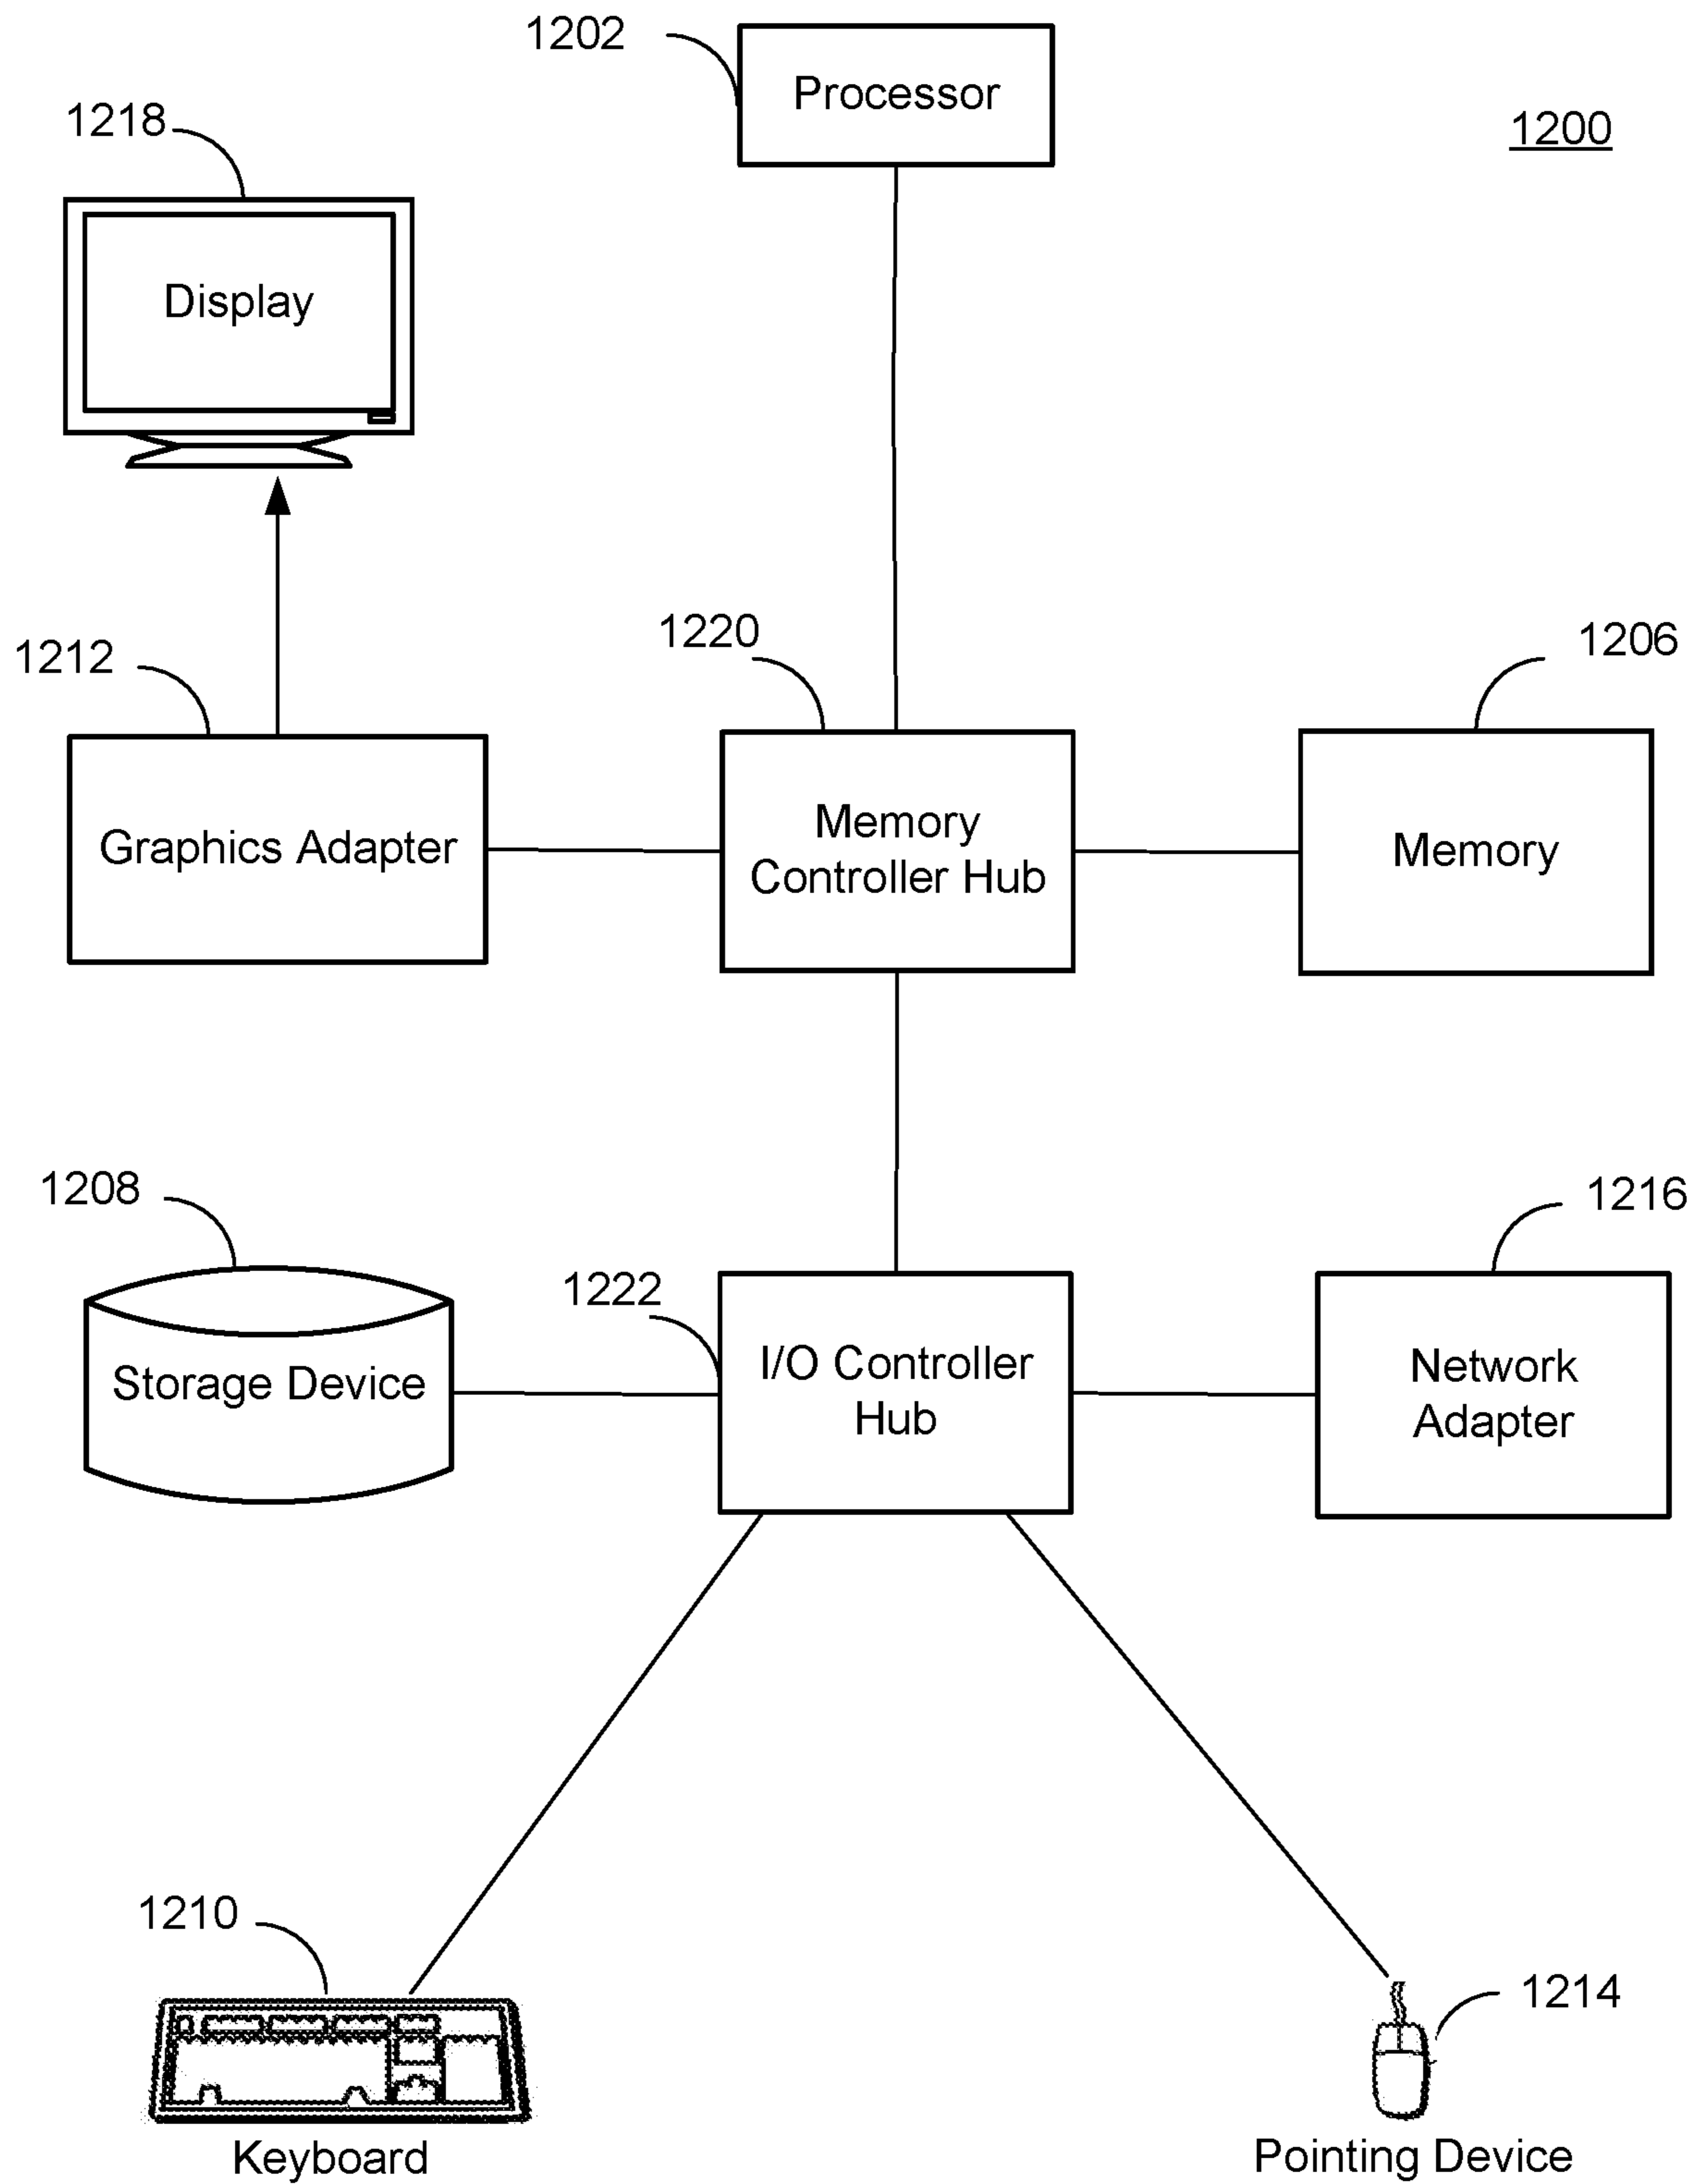


FIG. 12

## POLICY-DRIVEN DATA MANIPULATION IN TIME-SERIES DATABASE SYSTEMS

### CROSS-REFERENCE TO RELATED APPLICATION

This application is a continuation of the U.S. patent application Ser. No. 15/907,105, filed on Feb. 27, 2018, which claims the benefit of U.S. Provisional Application No. 62/464,289, filed on Feb. 27, 2017, each of which is incorporated by reference in its entirety.

### BACKGROUND

This disclosure relates generally to efficiently storing and processing data in a database system, and in particular to storing and processing time series data in a partitioned database system.

Time-series data is generated and processed in several contexts: monitoring and developer operations (DevOps), sensor data and the Internet of Things (IoT), computer and hardware monitoring, fitness and health monitoring, environmental and farming data, manufacturing and industrial control system data, financial data, logistics data, application usage data, and so on. Often this data is high in volume, for example, individual data sources may generate high rates of data, or many different sources may contribute data. Furthermore, this data is complex in nature, for example, a source may provide multiple measurements and labels associated with a single time. The volume of this stored data often increases over time as data is continually collected. Analytical systems typically query this data to analyze the past, present, and future behavior of entities associated with the data. This analysis may be performed for varied reasons, including examining historical trends, monitoring current performance, identifying the root cause of current problems, and anticipating future problems such as for predictive maintenance. As a result, operators are not inclined to delete this potentially valuable data.

Conventional systems fail to support the high write rates that are typical of many of these applications, which span across industries. For example, in Internet of Things (IoT) settings including industrial, agricultural, consumer, urban, or facilities, high write rates result from large numbers of devices coupled with modest to high write rates per device. In logistics settings, both planning data and actuals comprise time series that can be associated with each tracked object. Monitoring applications, such as in development and operations, may track many metrics per system component. Many forms of financial applications, such as those based on stock or option market ticker data, also rely on time-series data. All these applications require a database that can scale to a high ingest rate.

Further, these applications often query their data in complex and arbitrary ways, beyond simply fetching or aggregating a single metric across a particular time period. Such query patterns may involve rich predicates (e.g., complex conjunctions in a WHERE clause), aggregations, statistical functions, windowed operations, JOINS against relational data, subqueries, common table expressions (CTEs), and so forth. Yet these queries need to be executed efficiently.

Therefore, storing time-series data demands both scale and efficient complex queries. Conventional techniques fail to achieve both of these properties in a single system. Users have typically been faced with the trade-off between the horizontal scalability of “NoSQL” databases versus the query power of relational database management systems

(RDBMS). Existing solutions for time-series data require users to choose between either scalability or rich query support.

Traditional relational database systems that support database query languages such as SQL (structured query language) have difficulty handling high ingest rates: They have poor write performance for large tables, and this problem only becomes worse over time as data volume grows linearly in time. Further, any data deletion requires expensive “vacuuming” operations to defragment the disk storage associated with such tables. Also, out-of-the-box open-source solutions for scaling-out RDBMS across many servers are still lacking.

Existing NoSQL databases are typically key-value or column-oriented databases. These databases often lack a rich query language or secondary index support, however, and suffer high latency on complex queries. Further, they often lack the ability to join data between multiple tables, and lack the reliability, tooling, and ecosystem of more widely-used traditional RDBMS systems.

Distributed block or file systems avoid the need to pre-define data models or schemas, and easily scale by adding more servers. However, they pay the cost for their use of simple storage models at query time, lacking the highly structured indexes needed for fast and resource-efficient queries.

Conventional techniques that also fail to support an existing, widely-used query language such as SQL and instead create a new query language, require both new training by developers and analysts, as well as new customer interfaces or connectors to integrate with other systems.

### SUMMARY

The above and other issues are addressed by a computer-implemented method, computer system, and computer readable storage medium for dynamically creating chunks for storing records being added to a hypertable representing a partitioned database table. Embodiments of the method comprise receiving an insert request by a database system and processing the insert request. The insert request identifies a hypertable and one or more input records for inserting in the hypertable. Each record has a plurality of attributes including a set of dimension attributes that include a time attribute. The hypertable is partitioned into a plurality of chunks based on the dimension attributes. A chunk is specified using a set of values for each dimension attribute. For each record stored in the chunk, the value of each dimension attribute maps to a value from the set of values for that dimension attribute. A determination is made whether an input record should be stored in a new chunk or an existing chunk. For each new chunk being created, sets of values corresponding to each dimension attribute are determined and the new chunk is created for storing the input record. The hypertable is updated by storing the input record in the new chunk. The data stored in the updated hypertable is processed in response to subsequent queries that identify the hypertable.

Embodiments of a computer readable storage medium store instructions for performing the steps of the above method. Embodiments of a computer system comprise one or more computer processors and a computer readable storage medium storing instructions for performing the steps of the above method.

The features and advantages described in this summary and the following detailed description are not all-inclusive. Many additional features and advantages will be apparent to



one of ordinary skill in the art in view of the drawings, specification, and claims hereof.

### BRIEF DESCRIPTION OF THE DRAWINGS

The teachings of the embodiments can be readily understood by considering the following detailed description in conjunction with the accompanying drawings.

FIG. 1 is a block diagram of a system environment in which the database system operates, in accordance with an embodiment.

FIG. 2 illustrates partitioning of data of a database table, in accordance with an embodiment.

FIG. 3 shows processing of queries in a database system comprising a plurality of database nodes, in accordance with an embodiment.

FIG. 4A shows the system architecture of a query processor, in accordance with an embodiment.

FIG. 4B shows the system architecture of a chunk management module, in accordance with an embodiment.

FIG. 5 illustrates the process of inserting records into a hypertable stored across a plurality of database system nodes, in accordance with an embodiment.

FIG. 6 is a flowchart of the process of executing a query for processing records stored in a hypertable, in accordance with an embodiment.

FIGS. 7(A-C) illustrate partitioning of data of a database table to adapt to addition of locations to the database system according to an embodiment of the invention.

FIG. 8 shows a flowchart illustrating the process of modifying a data partitioning policy of a database system in response to addition of new locations to the database system, in accordance with an embodiment.

FIG. 9 illustrates selection of partitioning policy for creating chunks based on time attribute of the record, according to an embodiment.

FIG. 10 shows a flowchart of the process for selection of partitioning policy for creating chunks based on time attribute of the record, according to an embodiment.

FIG. 11 illustrates selection of partitioning policy for creating chunks based on time of receipt of a record by the database system, according to an embodiment.

FIG. 12 shows an architecture of a computer that may be used for implementing a database system node, in accordance with an embodiment.

### DETAILED DESCRIPTION

Embodiments of the invention include a database system that supports a standard query language like SQL and exposes an interface based on a hypertable that partitions the underlying data across servers and/or storage devices. The database system allows users to interact with data as if it were stored in a conventional database table, hiding the complexity of any data partitioning and query optimization from the user. Embodiments of the database system make a query language like SQL scalable for time-series data. The database system combines the best features of both RDBMS and NoSQL databases: a clustered scale-up and scale-out architecture and rich support for complex queries. Scaling up corresponds to running on larger individual servers, for example, machines with high numbers of CPUs or cores, or servers with greater RAM and disk capacity. Scaling up also includes increasing storage capacity of an existing database system by adding additional storage devices. Scaling out comprises increasing storage capacity of the database system by adding additional servers, for example, by sharding

the dataset over multiple servers, as well as supporting parallel and/or concurrent requests across the multiple servers.

### System Environment

FIG. 1 is a block diagram of a system environment in which the database system operates, in accordance with an embodiment. The system environment comprises a database system 110, one or more client devices 120, and a network 115.

The database system 110 comprises a query processor 130, a metadata store 140, and a data store 145. The database system 110 may include other components, for example, as illustrated in FIG. 2. The database system 110 receives database queries, for example, queries specified using SQL and processes them. The database system 110 may support standard SQL features as well as new user-defined functions, SQL extensions, or even non-SQL query languages such as declarative programming languages, a REST interface (e.g., through HTTP), or others.

The data store 145 stores data as tuples (also referred to as records) that may be stored as rows of data, with each row comprising a set of attributes. These attributes typically have a name associated with them (e.g., “time”, “device\_id”, “location”, “temperature”, “error\_code”) and a type (e.g., string, integer, float, boolean, array, json, jsonb (binary json), blob, geo-spatial, etc.), although this is not necessary in all cases. Attributes may also be referred to herein using the terms “fields”, “columns” or “keys”.

The data store 145 may store records in a standard database table that stores data in one or more files using conventional techniques used by relational database systems. The data store 145 may also store data in a partitioned database table referred to as a hypertable. A hypertable is a partitioned database table that provides an interface of a single continuous table—represented by a virtual view—such that a requestor can query it via a database query language such as SQL. A hypertable may be defined with a standard schema with attributes (or fields or column) names and types, with at least a time attribute specifying a time value. The hypertable is partitioned along a set of dimension attributes including the time attributes and zero or more other dimension attributes (sometimes referred to as the hypertable’s “space” attributes). These dimension attributes on which the hypertable is partitioned are also referred to as “partitioning key(s)”, “partition key(s)”, or “partitioning fields.” A hypertable may be created using a standard SQL command for creating a database table. Furthermore, queries to the hypertable may be made using database queries, for example, SQL queries.

The database system splits the hypertable into chunks. Each chunk stores a subset of records of the hypertable. A chunk may also be referred to herein as a data chunk or a partition. The database system 110 may distribute chunks of a hypertable across a set of one or more locations. A location may represent a storage medium for storing data or a system that comprises a storage medium for storing data, for example, a server. The storage medium may be a storage device, for example, a disk. The database system 110 may store data on multiple storage devices attached to the same server or on multiple servers, each server attached with one or more storage devices for storing chunks. A storage device may be attached to a remote server, for example, in a cloud-based system and a server of the database system provided access to the remote storage device for storing chunks.

The database system can store multiple hypertables, each with different schemas. Chunks within the same hypertable

often have the same schema, but may also have different schemas. The database system may also include standard database tables, i.e., traditional non-partitioned tables stored in the same database. Operations are performed against any of these tables, including multiple tables in a single query. For example, this can involve a SELECT that JOINS data between a hypertable and a standard non-partitioned table, or between two hypertables, or any more complex combination thereof. Or, it may involve inserting data into a hypertable and a standard non-partitioned table, or between two hypertables, or more complex combinations, as a single transaction.

In some embodiments, the database system **110** is comprised of one or more database system nodes (also referred to as database servers or just servers) that are connected over a network. Each node may include the same or similar components from FIG. 1, such as a query processor **130**, metadata store **140**, and data store **145**. The details of a database system node are described in FIG. 2. The metadata store **140** stores metadata describing the data stored in the data store **145** including descriptions of various hypertables and standard non-partitioned tables. The description includes various attributes of each table, the description of various chunks of a hypertable, and so on. The query processor **130** receives and processes queries as further described herein.

The database system **110** may be connected to requesters issuing database queries to the database system **110**. A requestor may be any source of the database queries, for example, a client device **120**, a webserver, application server, user workstation, or a server or machine that is sending the query on behalf on another origin (e.g., an intermediate server or middleware layer acting as a queue, buffer, or router such as for INSERTS, or an application acting on behalf of another system or user).

This connection from the requester often occurs over the network **115**, although it can also be on the same server executing the database system. For example, the network **115** enables communications between the client device **120** or any other requestor and the database system **110**. In one embodiment, the network uses standard communications technologies and/or protocols. The data exchanged over the network can be represented using technologies and/or formats including the open database connectivity (ODBC) format, the Java database connectivity (JDBC) format, the PostgreSQL foreign data wrapper (FDW) format, the PostgreSQL dblink format, the external data representation (XDR) format, the Google Protocol Buffer (protobuf) format, the Apache Avro format, the hypertext markup language (HTML), the extensible markup language (XML), Javascript object notation (JSON), etc.

The client device **120** can be a personal computer (PC), a desktop computer, a laptop computer, a notebook, a tablet PC executing an operating system. In another embodiment, the client device **120** can be any device having computer functionality, such as a personal digital assistant (PDA), mobile telephone, smartphone, wearable device, etc. The client device can also be a server or workstation, including running in a backoffice environment, within an enterprise datacenter, or within a virtualized cloud datacenter. The client device executes a client application for interacting with the database system **110**, for example, a browser **125**, a database shell, a web service application (such as .NET, Django, Ruby-on-Rails, Hibernate), a message broker (such as Apache Kafka or RabbitMQ), a visualization application, and so forth.

FIG. 1 and the other figures use like reference numerals to identify like elements. A letter after a reference numeral, such as “**120A**,” indicates that the text refers specifically to the element having that particular reference numeral. A reference numeral in the text without a following letter, such as “**120**,” refers to any or all of the elements in the figures bearing that reference numeral (e.g. “**120**” in the text refers to reference numerals “**120A**” and/or “**120N**” in the figures).

FIG. 2 illustrates partitioning of data as chunks for a hypertable, in accordance with an embodiment. Each of these chunks correspond to a portion of the entire dataset organized according to some partitioning function involving one or more of a record’s attributes. The attributes of the record that are used for partitioning the hypertable as chunks are referred to as dimension attributes. Accordingly, a chunk corresponds to an “n-dimensional” split of the hypertable (for  $n \geq 1$ ).

The database system **110** may implement a chunk as a file. In one embodiment, each chunk is implemented using a standard database table that is automatically placed on one of the locations (e.g., storage devices) of one of the database nodes (or replicated between multiple locations or nodes), although this detail may not be observable to users. In other embodiments, the placement of chunks on locations and/or database nodes is specified by commands or policies given by database administrators or users.

One of the dimension attributes is a time attribute that stores time-related values. The time attribute can be any data that can be comparable (i.e., has a  $>$  and  $\geq$  operator), such that data can be ordered according to this comparison function. Further, new records are typically associated with a higher time attribute, such that this value is commonly increasing for new records. Note that this value can be specified in the data record, and need not (and often does not) correspond to when data is inserted into the database. The following values may be used as a time attribute: datetime timestamps (including with or without timezone information), UNIX timestamps (in seconds, microseconds, nanoseconds, etc.), sequence numbers, and so on. In an embodiment, the hypertable is also split along a dimension attribute that represents a distinct identifier for objects or entities described in the database table (e.g., a device id that identifies devices, a server id that identifies servers, the ticker symbol of a financial security, etc.).

A chunk is associated with a set of values corresponding to each dimension attribute. For example, a hypertable may have two dimension attributes **d1** and **d2**. For a given chunk **C1**, the dimension attribute **d1** is associated with a set of values **S1** and the dimension attribute **d2** is associated with a set of values **S2**. Accordingly, each record stored in the chunk **C1** has a dimension attribute value that maps to a value in the set of values corresponding to the dimension attribute. For example, assume that a hypertable includes attributes time, device, and temperature. Also assume that time is a dimension attribute and a chunk is associated with a range of time [0:00:00-11:59:59.999]. If an input record has values {time: “1:00:00”, device: “A”, temperature: 65}, the chunk may store the input record since the value of the time dimension “1:00:00” falls within the range associated with the chunk, i.e., [0:00:00-11:59:59.999].

A set of values corresponding to a dimension attribute may represent a range of values but is not limited to ranges. For example, the set of values may represent a plurality of ranges that are not contiguous. Alternatively, the set of values may be specified by enumerating one or more values. For example, a dimension attribute **c1** may represent colors (e.g., “red”, “blue”, “green”, “yellow”), and a chunk may

store records that have the value of dimension attribute **c1** from the set {"red", "blue"} and another chunk may store records that have the value of dimension attribute **c1** from the set {"green", "yellow"}.

A given value of a dimension attribute may map to a value in the set of values corresponding to that dimension if the given value is identical to a value in the set of values. Alternatively, a given value **v1** of a dimension attribute may map to a value **v2** in the set of values corresponding to that dimension if the value **v2** is obtained by applying a transformation (for example, a function) to the given value **v1**. For example, database system **110** may use a hash partitioning strategy where the set of values corresponding to a dimension is specified as a range/set of values obtained by applying a hash function to the dimension attribute values. Accordingly, if a dimension attribute value is represented as **vx**, and **H** represents a hash function, a chunk **Cx** may be associated with a range **R[x1, x2]** (or set) of values for **H(vx)**. Accordingly, the chunk may store a record with dimension attribute value **v1** if **H(v1)** lies in the range **[x1, x2]**.

In an embodiment, the set of values may correspond to a plurality of dimension attributes. For example, the hash function specified in the above example may receive two or more inputs, each corresponding to a distinct dimension attribute, i.e., **H(v1, v2, . . .)**. Accordingly, a dimension of a chunk may be defined as a composite attribute comprising a plurality of dimension attributes of the hypertable.

FIG. 2 shows a hypertable **160** split into a plurality of chunks **210** along two dimension attributes, a time attribute and another dimension attribute referred to as the space attribute. In this example, each chunk is associated with a time range comprising a start time and an end time, and a space range comprising a contiguous range of alphabetical characters. For example, chunk **210a** stores a set of records that have the value of time attribute within the range **[0, 6]** and the value of space attribute within the range **[A, I]**. Similarly, chunk **210b** stores a set of records that have the value of time attribute within the range **[0, 6]** and the value of space attribute within the range **[J, R]**, and so on.

Different types of queries can be made to a hypertable, including those that only read from the hypertable (e.g., database **SELECT** statements), as well as those that modify the hypertable (e.g., database **INSERT**, **UPDATE**, **UPSERT**, and **DELETE** statements). Writes are typically sent to the chunks comprised of the latest time interval (but do not need to be), while queries may slice across multiple dimension attributes, for example, both time and space.

Although hypertables and chunks are referred to herein as tables, this term is not meant to be limiting, and a chunk could refer to a number of storage representations, including a traditional relational database table, a virtual database view, a materialized database view, a set of structured markup language (e.g., **XML**), a set of serialized structured data (e.g., **JSON**, **Google Protocol Buffers**, **Apache Avro**, **Apache Parquet**), or flat files (e.g., with comma- or tab-separated values).

#### Distributed Execution of Queries

FIG. 3 shows processing of queries in a database system comprising a plurality of database nodes, in accordance with an embodiment. A database system node **310a** receives database queries and may send one or more queries to chunks (that may be implemented as physical tables of the data), which are stored on the coordinator database system node or on other database system nodes. A database system node does not issue a query to a chunk if it determines that the chunk is not needed to satisfy the query. This determi-

nation uses additional metadata about the chunk, which may be maintained separate from or along with each chunk's data. Each database system node can also maintain additional metadata to allow it to efficiently determine a chunk's time interval and partitioning field's keyspace. The database system may maintain the metadata separate from or along with a chunk.

As shown in FIG. 3, the database system node **310a** receives a first database query **320**. The database system node **310a** determines that the data required for processing the received database query is on one or more database system nodes **310a**, **310b**, and **310c**. The database system node **310a** further sends queries **325a** and **325b** for processing the first query to the database system nodes **310b** and **310c**, respectively. All three database system nodes **310a**, **310b**, and **310c** process their respective queries using one or more chunks of data stored on their respective nodes. In the example, illustrated in FIG. 3, if the database system node **310a** determines that the data required for processing the first query is stored only on database system nodes **310a** and **310b** but not on **310c**, the database system node **310a** sends a query for processing to **310b** but not to **310c**. In other embodiments of the system, the queries **325a** and **325b** sent to the other nodes **310b** and **310c** are the same as the first query **320**, and the queries or requests sent to the other nodes can be in a different query language, format, or communication protocol as the first query. In some embodiments of the system, the queries **325a** and **325b** maybe identical to each other, while in others they are different. Further, in other embodiments, node **310a** does not store chunks itself, but only processes the query **320** and issues the corresponding queries **325** to other database nodes.

The database system node **310a** that receives the database query may determine that the query to the hypertable does not involve a particular chunk's data—for example, because the query specified a time period different than that associated with the chunk, or if the query specifies a dimension attribute (e.g., an IP address, device ID, or some location name) that is associated with a different chunk. In this situation, the first database system node does not issue a query to this particular chunk (which may be located on itself or on a different node). This determination by both the first database system node and any other database system nodes may be performed by the query processor **130** present on each database system node that processes queries.

Any database system node may receive a query from a requester and the query processor **130** running on this database system node determines how to plan and execute the query across the entire cluster of one or more nodes. This database system node sends a query (a "subquery") to zero or more other nodes in the system. Subsequently, the database system node(s) that receive a subquery from the first database system node include a query processor **130** that determines how to plan and execute the query locally.

In an embodiment, this process is extended to additional levels of subqueries and involved planners. In an embodiment, the database system performs this partitioning in a recursive fashion. For example, the chunk that is being stored on one of the nodes could itself be further partitioned in time and/or by an additional partitioning key (either the same or different than the partitioning key at a higher level), which itself could be distributed among the node (e.g., on different disks) or even to other nodes. In such a scenario, a chunk can act as another hypertable.

In some embodiment, the database system performs the query processing using only the query processor **130** on the first database system node. Accordingly, the complete query

plan is generated by the first node and sent to nodes that are determined to store chunks processed by the query. The remaining nodes that receive the query plan (or some portion thereof) simply execute the received query plan without having to generate a portion of the query plan. In other embodiments, the database system implements less homogeneous functionality across nodes, such that a first set of one or more nodes receives queries and plans and executes the queries against a second disparate set of one or more nodes that store the chunks.

#### System Architecture

FIG. 4A shows the system architecture of a query processor, in accordance with an embodiment. The query processor 130 comprises components including a connector 410, a query parser 415, a query planner 425, a query optimizer 430, an execution engine 435, and a query plan store 455. A query processor 130 receives a query in some query language, such as SQL, which specifies the tables or datasets on which the query will apply (i.e., read or write data). A query or database query may represent a request to read data (e.g., SELECT statements in SQL) or modify data (e.g., INSERT, UPDATE, and DELETE statements in SQL) from the database.

The query parser receives this request and translates it to a query representation that is easier to process. For example, the query parser 415 may generate a data structure representing the query that provides access to the information specified in the query. The query optimizer 430 performs transformation of the query, for example, by rewriting portions of the query to improve the execution of the query. The query planner takes this machine-readable representation of the query, which is typically declarative in nature, and generates a plan specifying how the query should be executed against the stored data, which may be stored in memory (e.g., RAM, PCM) and/or on some type of non-volatile storage media (e.g., flash SSD, HDD). The query processor 130 stores the generated plan in the query plan store 455. The execution engine 435 executes the query against the stored data, and returns the results to the requester. The connector 410 allows the query processor 130 to connect to remote systems, for example, to access data stored in remote systems.

FIG. 4B shows the system architecture of a chunk management module, in accordance with an embodiment. The chunk management module 170 further comprises a chunk selection module 445, and a chunk creation module 450. The chunk selection module 445 implements a chunk selection function that determines a chunk for storing a given record. The chunk selection module 445 determines whether an existing chunk can be used for storing the record. If the chunk selection module 445 determines that none of the existing chunks can be used for storing the record, the chunk selection module 445 determines that a new chunk needs to be created and invokes the chunk creation module 450 for creating the chunk. If the chunk selection module 445 determines that a new chunk needs to be created, the chunk selection module 445 determines various parameters describing the new chunk. For example, the chunk creation module 450 determines the sets of values corresponding to different dimension attributes of the records that define the chunk boundaries. Accordingly, records stored in the chunk have dimension attribute values such that each dimension attribute has a value that maps to a value in the set of values corresponding to the chunk's dimension. For example, if a chunk has two dimension attributes, based on time and a device id, then each record stored in the chunk has a time attribute that falls within the chunk's time range and a device

id from the set of device ids associated with the chunk. The chunk creation module 450 determines a location for creating the chunk and creates the chunk on the location.

The database system 110 stores in the metadata store 140, metadata 155 describing the chunk. The metadata for a chunk includes information associating the chunk with the hypertable. Other type of metadata describing the chunk includes a name of the chunk, the various sets of values of the dimension attributes (for example, time ranges for the time attribute, and so on), information describing constraints and indexes for the chunk, and so on. The database system 110 may store other metadata associated with the chunk, e.g., access statistics and data distribution statistics to aid query planning.

A hypertable may be associated with certain policy configurations, for example, indexes, constraints, storage parameters (e.g., fillfactor settings, parallel worker settings, autovacuum settings, etc.), foreign key relationships, and so on. In an embodiment, each chunk of the hypertable implements the policy configurations of the hypertable containing the chunk. Accordingly, when creating a chunk, the chunk creation module 450 may also create structures such as indexes for the chunk and update metadata to specify constraints, foreign key relationships, and any other policy configurations for the chunk. Examples of constraints defined for a chunk include UNIQUE, NOT NULL, CHECK CONSTRAINT (i.e., timestamp between range), FOREIGN KEY, and EXCLUSION constraints. The chunk management module 170 continues to manage the chunk once it is created, for example, by reindexing old chunks periodically, moving old chunks to slower storage devices over time, adding secondary constraints through dynamic inspection, and so on.

In an embodiment, the chunk management module 170 monitors the sizes of the chunks that were recently created. A recently created chunk (or a recent chunk) refers to a chunk that was created within a threshold time interval of the current time. The size of the threshold time interval may be configurable. The size represents the amount of data that is stored in the chunk, for example, the chunk's size of bytes, its number of rows, and so on. The chunk management module 170 adjusts sets of values of the dimension attributes for new chunks being created based on the size of the recently created chunks. Accordingly, if the chunk management module 170 determines that one or more recently created chunks store data that exceeds certain high threshold values, the chunk management module 170 adjusts the sets of values of one or more dimensions so that they have fewer elements than the corresponding sets of values of the recently created chunks. For example, if the chunk management module 170 determines that the recently created chunks had a range of 12 hours for the time attribute, the chunk management module 170 may decrease the range of time attributes of new chunks being created to be 10 hours. Alternatively, if the chunk management module 170 determines that one or more recently created chunks store data that is below certain low threshold values, the chunk management module 170 adjusts the sets of values of one or more dimensions so that they have more elements than the corresponding sets of values of the recently created chunks that were below the low size thresholds. For example, if the chunk management module 170 determines that the recently created chunks had a range of 12 hours for the time attribute and stored very few records, the chunk management module 170 may increase the range of time attributes of new chunks being created to be 15 hours.

In an embodiment, the chunk management module 170 monitors one or more performance metrics for the chunks that were recently created. The chunk management module 170 adjusts the sets of values of dimension attributes for new chunks being created based on the performance metrics for the chunks that were recently created. For example, the chunk management module 170 may monitor insert rate and query execution time. For example, if the chunk management module 170 determines that for the current sizes of chunks the insert rate of records has fallen significantly (e.g., since the database system has started swapping to disk), then the chunk management module 170 determines the sets of values of dimension attributes of new chunks being created such that the new chunks are smaller.

In an embodiment, chunk management module 170 keeps statistics describing chunks processed by each distinct query, for example, the number of chunks processed by each query. The chunk management module 170 uses this statistical information to determine sets of values for dimension attributes of new chunks being created so as to improve performance. In an embodiment, the chunk management module 170 monitors the dimension attribute boundaries specified in queries. If the chunk management module 170 determines that commonly received queries have certain pattern of boundaries, for example, a pattern of time alignment (e.g., typical queries request data for a day between midnight and midnight), then the chunk management module 170 aligns newly created chunks to match these boundaries. As another example, if the current chunks have one hour time attribute ranges and the chunk management module 170 determines that the queries are typically accessing data at an interval of a size of a full day, the chunk management module 170 increases the chunk sizes to reach a size more aligned with the access patterns, yet one that still retains a high insert rate. For example, the chunk management module 170 may increase the time attribute range to be 12 hours, e.g., if 12 hours gives a higher insert rate compared to a 24-hour range.

In an embodiment, the chunk management module 170 determines the sets of values of the dimension attributes of chunks being created based on ranges of dimension attributes specified in queries received by the database system. For example, if the chunk management module 170 is creating chunks with time attribute ranges from 11 pm to 11 pm, and the chunk management module 170 determines that the queries received are accessing data from midnight to midnight, the chunk management module 170 shifts the time range of the chunks being created to match the time ranges of the queries. This improves the performance of queries by avoiding the need to unnecessarily scan two chunks rather than one.

In an embodiment, the chunk management module 170 distributes chunks across a plurality of locations based on the properties of the storage media of each location. The chunk management module 170 identifies the storage medium for storing the new chunk and accesses properties of the storage medium, for example, properties describing a rate of access of data stored on the storage medium. The chunk management module 170 determines a number of chunks from the plurality of chunks being assigned to a location based on the properties of the storage medium corresponding to that location. For example, the chunk management module 170 accesses metrics describing the rate at which a storage medium accesses random data. Certain storage mediums, e.g., solid-state drives (SSDs) and random-access memory (RAM), can handle random reads much better than spinning hard disk drives (HDDs). Accord-

ingly, the chunk management module 170 assigns more chunks from the plurality of chunks to a location having a storage medium with faster access time for random accesses.

In one embodiment, the chunk creation module 450 creates a new chunk—and “closes” an existing one—when the existing chunk approaches or exceeds some threshold size (e.g., in bytes on disk or in memory, in its number of rows, etc.). Each chunk is represented by a start and end time (defining its interval). With a purely size-based approach, however, the database system would not know a priori the end time of a newly-created chunk. Thus, when a chunk is first created, the chunk’s end time is unset; any row having time greater than (or equal to) the start time is associated with the chunk. However, when a chunk’s size approaches or exceeds some threshold, the query planner 425 closes the chunk by specifying its end time, and the chunk creation module 450 creates a new chunk. This new chunk starts at the time the old chunk ends. With this approach, the chunk has an indeterminate end time for a chunk until it is closed.

A similar logic is applied to an indeterminate start-time. It is also possible for an initial chunk to have both an indeterminate start and end time. An embodiment of the database system performs this determination and chunk creation asynchronously or in the background, while another performs these actions during the process of inserting a (set of) row(s) from the received batch to the chunk. The creation of the new chunk at insert time can happen in a variety of ways: before inserting the rows (the query planner 425 decides that the existing chunk is too full already, and creates a new chunk to insert into); after inserting the rows into the chunk; or in the middle of inserting the rows (e.g., the query planner 425 decides the chunk only has space for a subset of the rows, so the subset is inserted into the current chunk and the remainder of the set is inserted into a newly created chunk).

In other embodiments, the database system defines a chunk as having a particular time interval (that is, both a start and end time) when the chunk is created. Then the system creates a new chunk when needed, e.g., when new data is to be inserted to a time interval that does not yet exist. In one embodiment, the database system also employs a maximum size even with this approach, so that, for example, a second chunk is created with the same time interval as the first chunk if the size is approached or exceeded on the first chunk, and the query planner 425 writes new data to only one of the chunks. Once a second chunk is created, the database system may rebalance data from the first to second chunk. In another embodiment, rather than overlap the time intervals of the first and second chunk, the first chunk’s end time is modified when the second chunk is created so that they remain disjoint and their time intervals can be strictly ordered. In another embodiment, the database system performs such changes asynchronously, so that an over-large chunk is split into a first and second chunk as a “background” task of the system. Further, in another embodiment, this second chunk is created when an insert occurs to a time value that is sufficiently close to the end of a chunk’s time range, rather than only when a record’s dimension attributes (e.g., time) fall outside the dimensions of any existing chunks. In general, many of the variations of the database system’s chunk management may be performed either synchronously at insert time or asynchronously as a background task. Size- and interval-based chunking is further described below.

In an embodiment, the chunk creation module 450 performs collision detection to ensure that the new chunk(s) have sets of dimension attributes that are disjoint from existing chunks. For example, assume that the chunk cre-

ation module is creating chunks with a time range spanning 24 hours. If the previous chunk stored data with time attribute values until midnight (exclusive) on a date January 1, the chunk creation module **450** next creates chunks with time attribute values from midnight (inclusive) on January 2 to the following midnight (exclusive). As another example, if the chunk creation module **450** is creating chunks with 18-hour intervals of time attribute, if the previously created chunk covered a time interval from midnight to 3 am, the chunk creation module **450** next creates a new 18-hour chunk spanning a time interval from 3 am to 9 pm for the time attribute. The chunk creation module **450** can create multiple chunks having the same time range but having different sets of values for other dimension attributes.

The chunk creation module **450** may adjust chunk boundaries based on various criteria, some of which may be conflicting. As an example, consider that the database system has one chunk with a time interval that ends at 3 am, and another chunk from noon to the following midnight. The database system may next receive a request to insert a record having a time attribute value of 4 am. Even if the chunk creation module **450** may be creating chunks with a time range spanning 12 hours, in this scenario, the chunk creation module **450** may create a new chunk spanning only a 9 hour time interval from 3 am to noon in order to enforce disjointness. In some embodiments, the chunk management module **170** determines after a chunk is created that the ranges (or set of values) of the chunk are likely to overlap other chunks created. In these embodiments, the chunk management module **170** modifies the existing ranges of the chunk to ensure that the ranges are disjoint from other chunks.

In some embodiments, across different partitions, the database system may align chunk start and end times or maintain them independently. In other words, the system may create and/or close all of a hypertable's chunks at the same time, or different partitions can be managed distinctly from one another. In other embodiments, there may be special overflow chunks where data that cannot be placed in some existing chunks is placed either temporarily or permanently.

The system architecture illustrated in these figures (for example, FIGS. 1-4) are meant to be illustrative; other embodiments may include additional or fewer components, some of these components might not always be present (e.g., a query parser or cache), or these components may be combined or divided in a variety of way (e.g., the query planner, query optimizer, and execution engine). It is understood that such a representation or division would not change the overall structure and function of the database system. For example, one would understand that the methods described herein could be implemented in a system that includes one component performing both query planning and executing, or in a system that includes a separate component for planning, which then passes the plan to an executor engine for execution.

#### Inserting Data in a Hypertable

FIG. 5 illustrates the process of inserting records into a hypertable stored across a plurality of database system nodes, in accordance with an embodiment. The database system **110** receives **510** an insert query (which we also call an insert request). The insert query identifies a database table, for example, a hypertable, chunk, or a standard non-partitioned database table and specifies one or more records to be inserted into the database table. The database system **110** may store records as a hypertable comprising a plurality of chunks, each chunk stored in a distinct location.

Upon receiving **510** the insert query, the query parser **415** parses the insert query. The query planner **425** processes the query, and determines if the query specifies a hypertable, chunk, or a standard non-partitioned database table. If the insert query specifies a standard database table or a chunk, the query planner **425** executes the insert on the specified chunk or the standard database table in conjunction with the execution engine **435** and returns the result(s).

If the query specifies a hypertable, the query processor **130** performs the following steps for each record specified in the insert request. The query processor **130** identifies the values of the dimension attributes in the input record. The query processor **130** determines whether the input record should be stored in an existing chunk or in a new chunk that needs to be created. In an embodiment, the query processor **130** determines whether the one or more dimension values of the input record map to values from the set of dimension attribute values of existing chunks storing data of the hypertable; this determination is made to decide whether the record can be stored in an existing chunk.

In an embodiment, the query processor **130** provides **520** the dimension attributes as input to a selection function of the chunk selection module **445** that determines whether the record should be stored in an existing chunk or whether a new chunk needs to be created for storing the record. If the selection function finds an existing chunk that matches the record, the selection function outputs information identifying the existing chunk. If the selection function determines that none of the existing chunks can be used to store the record, the selection function outputs a value (for example, a negative number) indicating that a new chunk needs to be created. The chunk creation module **450** determines **540** based on the output of the selection function, if the record matches an existing chunk. If the chunk creation module **450** determines **540** that the record matches an existing chunk, the chunk selection module **445** also identifies the location of the existing chunk, for example, whether the existing chunk is local (i.e., on the current database system node) or remote (i.e., on another database system node). This location can specify a location explicitly or implicitly, including specifying a name of a local database table, the name of a remote database table, the name or network address or a remote server, and so on. Accordingly, the query processor **130** inserts **550** the record in the existing chunk.

If the chunk creation module **450** determines **540** based on the output of the selection function that a new chunk needs to be created for storing the record, the chunk creation module **450** determines **560** a configuration of the new chunk comprising sets of values corresponding to different dimension attributes for the new chunk. The chunk creation module **450** may further identify a location for creating the new chunk (including identifying a specific storage device or instead identifying a specific database system node, wherein the identified node in turn identifies a specific storage device attached to it). The chunk creation module **450** creates a new chunk based on the configuration of the new chunk and the identified location. The query processor **130** inserts **580** the record in the new chunk that is created.

The chunk selection module **445** may determine that a record cannot be inserted in an existing chunk based on various criteria. A record cannot be inserted in any existing chunk if the dimension attributes of the record do not match the configurations of any existing chunks. In some embodiments, even if the dimension attributes of the record match the configuration of an existing chunk, the chunk selection module **445** may determine that the record cannot be inserted into the chunk based on certain policy consider-

ations. For example, the chunk selection module **445** may determine that the existing chunk is storing more than a threshold amount of data and no new records should be added to the chunk. Accordingly, the chunk selection module **445** determines that the record cannot be added to the existing chunk and the database system cannot insert the record in any existing chunk.

To create a chunk locally or to insert the record in a chunk stored locally, i.e., on the current database system node executing the above steps, the database system may perform a function call. To create a chunk remotely or to insert the record in a chunk stored remotely, i.e., on a database system node different from the current database system node, the database system may perform a remote call, for example, a remote procedure call (RPC) or a remote SQL query execution. The instructions executed for creating a chunk or inserting a record into a chunk may also depend on the location of the chunk, for example, the type of storage medium used for storing the chunk.

Although FIG. 5 describes the steps in terms of a selection function, other embodiments can use different functions to compute different values, for example, a first function to determine whether the record should be stored in an existing chunk and a second function to describe a new chunk if the first function determines that the record cannot be stored in any existing chunk.

If multiple chunks reside on the same location, rather than using a separate message for each insert query, the query processor **130** may send multiple queries in a single message, or it may also send the multiple records to be inserted in a single query in a single message. If the chunks involved in an insert query reside on multiple nodes, in some embodiment the database system node contacts a query or transaction coordinator for additional information that is used and/or transmitted when subsequently communicating with other database nodes as part of the insert process.

In some embodiments, the query processor **130** handles a lack of a timely response or an error in a variety of ways. If a chunk is replicated between multiple nodes, or the record-to-chunk determination process results in more than one chunk, the query processor **130** issues an insert request to one or more of these chunks, discussed further. Finally, the query planner **425** collects any result(s) or status information from the insert queries, and returns some result(s) or status information to the requester.

In some embodiments, the database system **110** performs several steps to determine the chunk to which a record belongs, many of which involve using metadata. First, the query planner **425** determines the set of one or more partitions that belong to the hypertable at the time specified by the record (i.e., the value of the record's time attribute). If this partitioning is static, the query planner **425** uses metadata about the hypertable itself to determine this partitioning.

If this partitioning changes over time, the query planner **425** uses the record's time attribute to determine the set of partitions. In one embodiment, this determination involves first using the row's time attribute value to determine a particular epoch (time interval), then using this epoch to determine the set of partitions. This partitioning may change in the context of system reconfiguration (or elasticity) as described below. Second, the query planner **425** determines the partition (from amongst this set of one or more partitions) to which the record belongs, using the value(s) of the record's dimension attribute(s). For each of the dimension attributes used for partitioning in the hypertable, this step may involve applying some function to its value to generate

a second value. A variety of functions may be employed for this purpose, including hash functions (e.g., Murmur hashing, Pearson hashing, SHA, MD5, locality-sensitive hashing), the identity function (i.e., simply return the input), a lookup in some range-based data structure, or some other prefixing or calculation on the input. Third, using this second value (the function's output), the query planner **425** determines to which partition the second value belongs. For example, this step could involve a range lookup (e.g., find the partition [x, y] such that the second value is between x and y, inclusive and/or exclusive), a longest-prefix match on the partition (determine the partition that, when represented by some binary string, has the greatest number of most significant bits that are identical to those of the second value), taking the second value "mod" the number of nodes to determine the matching partition number, or the use of consistent hashing, among other matching algorithms. If the hypertable is partitioned using more than one key, then a function could be applied to more than one input (or functions could be separately applied to multiple inputs), leading to one or more second values (outputs) that would be used to determine the partition to which a record belongs. Finally, each partition for each dimension is associated to a set of chunks (i.e., those chunks which store this partition yet may differ in their time ranges); the query planner **425** then determines a chunk from this set based on the record's time attribute.

Other embodiments implement the step of determining the chunk to which a record belongs in alternate ways. For example, the database system skips the process of first determining a record's chunk based on its epoch, and instead first determines a set of chunks associated with the record's time. The query planner **425** computes a function on the record's partition key(s) to determine the second value(s), and compares this second value against the partition information associated with each chunk in order to select one. These processes can be implemented via a variety of data structures, including hash tables, linked lists, range trees, arrays, trees, tries, etc.

There are a variety of other optimized ways to implement the process by which the query planner **425** inserts a batch's data into chunks, without changing its basic functionality. For example, rather than performing all these steps for every record, the query planner **425** can cache information it determines during its per-record analysis, such as the hypertable's chunks for a given time or time period.

Other embodiments perform the steps for processing a batch in different ways. For example, after determining the first record's chunk, the query planner **425** scans through the rest of the batch, finding all other records associated with the same chunk (if any exist). The query planner **425** then inserts these records into the selected chunk, and deletes them from the batch. The query planner **425** then repeats this process: selecting a record in the (now smaller) batch, scanning the rest of the batch to find records with a similar chunk association, sending that set of one or more records to the second chunk, and then repeating this process until the batch is empty.

The insertion process above describes a record as being associated with a single chunk. Alternatively, a record could map to multiple chunks. For example, the chunking process might create more than one chunk during a particular interval (e.g., if the size of inserted data exceeds some threshold), as described herein, in which case the selection function chooses one, e.g., randomly, round robin, or based on their sizes. As another example, the database chooses to insert the record into multiple chunks to replicate data for

reliability or high availability. Such replication can be performed by the query planner **425** as part of the same steps described above, or the query planner **425** first inserts each of the records into a primary chunk, and then the database system **110** replicates the inserted record to the chunk's replica(s).

In an embodiment, the database system **110** replicates the chunks such that different chunks of the same hypertable may be stored with a different number of replicas. Furthermore, the database system may determine the number of replicas for a chunk based on the age of the chunk. For example, recent chunks may be replicated a greater number of times than older chunks. Furthermore, older chunks that have more than a threshold age may not be replicated. The database system **110** may determine the age of a chunk based on the values of the time attribute of the chunk. For example, a chunk that stores records having time attribute within a range [t1, t2] may be determined to be older than a chunk that stores records having time attribute within a range [t3, t4] if the time range [t1, t2] is older than the time range [t3, t4], for example, t2 < t3. Alternatively, the age of the chunk may be determined based on the time of creation of the chunk. For example, a chunk created a week ago has an age value that is greater than a chunk created today.

In an embodiment, the database system replicates different chunks to locations having different characteristics. The database system selects a location having particular characteristics based on the configuration of the chunk. For example, the database system stores and/or replicates recent chunks which are regularly being accessed (for inserts or selects) on fast storage media (e.g., SSDs), while the database system stores and/or replicates old chunks on slower storage media (e.g., HDDs).

In an embodiment, the database system reuses replication techniques that apply to the database's underlying tables, namely, physical replication of the entire database and cold/hot standbys, logical replication of individual tables, as well as backups. It also uses the database's write-ahead log (WAL) for consistent checkpointing. In other words, even though replication or backup policies are defined (or commands issued) on the hypertable, the system performs these actions by replicating or checkpointing the hypertable's constituent chunks. In another embodiment, replication and high availability is implemented directly by the database system by replicating writes to multiple chunk replicas (e.g., via a two-phase commit protocol), rather than by using the database's underlying log-based techniques.

In an embodiment, the database system allows different policies to be defined based on chunk boundaries, e.g., a higher replication level for recent chunks, or a lower replication level on older chunks in order to save disk space.

In an embodiment, the database system also moves chunks between locations when they age (e.g., from being stored on faster SSDs to slower HDDs, or from faster or larger servers to slower or smaller servers). The database system associates each hypertable with a threshold age value. The database system further associates locations with types. For example, different types of locations may have different access time, different storage capacity, different cost, and so on. If the database system identifies a chunk of the hypertable having an age value greater than the threshold age value of the hypertable, the database system moves the identified chunk from a location having a particular type to another location having a different type. As a result the database system may store different chunks of the same hypertable in different types of location. Furthermore, the database system automatically changes the mapping of the

chunks of the hypertable to locations over time as newer chunks are received and existing chunks get older. In another embodiment, this movement only happens when requested by a command (e.g., from an external process or database user), which specifies the age associated with the hypertable and the locations between which to move any selected chunks.

Processing Queries Reading Data

FIG. 6 is a flowchart of the process of executing a query for processing records stored in a hypertable, in accordance with an embodiment. The database system receives **610** a query for reading data (e.g., via a SELECT statement in SQL). Upon receiving a query, the query parser **415** parses the query (optionally using a cache of parsed queries). The query planner **425** processes the query and determines if any table specified in the query corresponds to a hypertable, chunk, or a standard non-partitioned database table. The database system performs the following steps in these different scenarios, each resulting in some result being returned to the requester (or some form of error if any problems occur).

For every hypertable specified in the first query, the query planner, in conjunction with the execution engine **435**, performs the following steps. First, the query planner **425** analyzes the query to determine **620** the set of chunks that may contribute results to the query's answer. This analysis typically involves the constraints specified by the query's predicates as well as metadata that the database system **110** maintains about chunks. For example, these constraints may be based on the value of a particular field (e.g., selected rows must have a device identifier that equals either 100 or 450), or they may include some type of time range (e.g., selected rows must specify that their time value is within the past hour, or between July 2016 and August 2016). Metadata stored about each chunk may specify, among other things, the range of time and any other partitioning key(s) associated with a particular chunk. For example, a chunk might be storing the last day of data for device identifiers between 0 and 200. These examples are simply illustrative and a variety of techniques that the system may employ are described herein. The query planner **425** uses the metadata to determine the appropriate chunks, e.g., a device identifier of 100 will be associated with the chunk storing device identifiers between 0 and 200.

The following steps **630**, **640**, **650**, and **660** are repeated for each chunk determined. The query planner **425** uses metadata to determine the location(s)—e.g., storage devices such as local or network-attached disk(s), or other database system node(s)—at which these chunk(s) are being stored. These chunks may be stored on a location local or remote to the query planner **435**. The query planner **425** determines **640** whether the chunk is stored locally or on a remote server. If the query planner **425** determines that the chunk is stored in a local location, the query planner **425** queries the local chunk (e.g., via direct function calls) or else the query planner **425** sends **660** a query to the remote location storing the chunk (e.g., by issuing SQL queries such as via foreign data wrappers, by sending remote procedure calls (RPCs), etc.). Furthermore, the query planner **425** may change the query execution or plan depending on the properties of the location that stores them (e.g., type of disk or node). When multiple chunks share the same location, the query planner **425** can generate a single query for the location's set of chunks or a separate query per chunk, and these separate queries can be sent in a single message to the location or as a separate message per query.



The query planner 425 issues queries to these locations and waits for their results. If some locations are not responding after some time or return errors, the query planner 425 can take several different options, including retrying a query to the same location, retrying a query to a different location that replicates the chunk, waiting indefinitely, returning a partial result to the client, or returning an error. The query planner 425 receives 670 or collects the results of these queries and merges the results. Depending on the query, the results, metadata, and additional information, the query planner 425 optionally may determine that it needs to query additional chunks to resolve the first query (e.g., when “walking back in time” from the latest time interval to older intervals in order to find some number of values matching a particular predicate).

Depending on the query, the query planner 425 may perform 680 post-processing of the results. Such post-processing includes taking a union over the returned results, performing an aggregation like a SUM or COUNT over the results, sorting the merged results by a specific field, taking a LIMIT that causes the system to only return some number of results, and so on. It may also involve more complex operations in merging the chunks’ results, e.g., when computing top-k calculations across the partial results from each chunk. Finally, the system returns the result(s) of this first query. The result of the query may comprise one or more tuples or an error code if the processing of the query resulted in an error.

In some embodiment, a query across multiple database nodes may also involve the use of a query or transaction coordinator, such that the coordination is contacted for additional information that is used and/or transmitted when subsequently communicating with other database nodes as part of the query process.

A node may also receive a query to a chunk or chunks, e.g., because it is the recipient of a query generated by the processing of the first query to a hypertable. For every chunk specified in the query, the query planner 425 performs the following steps. The query planner 425 plans and executes the query on the local chunk. This uses query planning techniques including choosing and optimizing the use of indexes, performing heap scans, and so forth. The query planner 425 receives the results of the query. Third, depending on the query, the query planner 425 may also post-process the results (e.g., sorting the data, performing an aggregation, taking the LIMIT, etc. as described above). It then returns the query’s result(s).

A database system node may receive a query to a traditional database table, which involves processing the query in a standard way: planning and executing the query on the specified table, receiving the results, post-processing the results optionally, and returning the result(s).

The query may also specify multiple tables or joins between tables. The database system’s processing depends on the types of tables specified (e.g., hypertables, chunks, standard non-partitioned tables) and is related to the steps above, although individual steps may differ or additional steps may be required based on the actual query.

#### Alternative Embodiments for Processing Queries Based on Hypertables

Ideally database users should be able to interact with time-series data as if it were in a simple continuous database table. However, for reasons discussed above, using a single table does not scale. Yet requiring users to manually partition their data exposes a host of complexities, e.g., forcing

users to constantly specify which partitions to query, how to compute JOINS between them, or how to properly size these tables as workloads change.

To avoid this management complexity while still scaling and supporting efficient queries, the database system hides its automated data partitioning and query optimizations behind its hypertable abstraction. Creating a hypertable and its corresponding schema is performed using simple SQL commands, and this hypertable is accessed as if it were a single table using standard SQL commands. Further, just like a normal database table, this schema can be altered via standard SQL commands; transparently to the user, the database system atomically modifies the schemas of all the underlying chunks that comprise a hypertable.

In an embodiment, the database system provides this functionality by hooking into the query planner of a relational database like PostgreSQL, so that it receives the native SQL parse tree. It uses this tree to determine which servers and hypertable chunks (native database tables) to access, how to perform distributed and parallel optimizations, etc.

Many of these same optimizations even apply to single-node deployments, where automatically splitting hypertables into chunks and related query optimizations still provides a number of performance benefits. This is especially true if the chunks are distributed across the various locations of a node (e.g., across multiple local or network-attached disks). In an embodiment, the placement of chunks on database nodes is specified by commands or policies given by database administrators or users.

In an embodiment, the database system partitions its hypertable in only a single dimension—by time—rather than two or more dimensions (for example, time and space dimensions). For example, partitioning based on a single time dimension may be used for deployments of the database system on a single node rather than a cluster of nodes.

Additionally, hypertables can be defined recursively. In particular, a hypertable’s chunk can be further partitioned (by the same or different partitioning key, and with the same or different time intervals) and thus act like another hypertable.

Chunks are dynamically created by the runtime and sized to optimize performance in both cluster and single-node environments. Partitioning a hypertable along additional dimension attributes (in addition to time) parallelizes inserts to recent time intervals. Similarly, query patterns often slice across time or space, so also result in performance improvements through chunk placements disclosed herein.

The placement of these chunks can also vary based on deployment, workload, or query needs. For example, chunks can be randomly or purposefully spread across locations to provide load balancing. Alternatively, chunks belonging to the same region of the partitioning field’s keyspace (for example, a range of values or hashed values, or a set of consecutive values of the key), yet varying by time intervals, could be collocated on the same servers. This avoids queries touching all servers when performing queries for a single object in space (e.g., a particular device), which could help reduce tail latency under higher query loads and enable efficient joins.

The database system determines where a chunk should be placed when it is created; this determination is based on a variety of one or more metrics, including performed randomly or via a round-robin distribution strategy, based on server load (e.g., request rate, CPU utilization, etc.), based on existing usage (e.g., size of existing chunks in bytes or number of rows), based on capacity (e.g., total memory or

storage capacity, free memory, available storage, number of disks, etc.), based on configured policy or specified by an administrator, and so forth. The database system or administrator may also choose to relocate (move) or replicate chunks between servers.

Even in single-node settings, chunking still improves performance over the vanilla use of a single database table for both read and write queries. Right-sized chunks ensure that most or all of a table's indexes (e.g., B-trees) can reside in memory during inserts to avoid thrashing while modifying arbitrary locations in those indexes. Further, by avoiding overly large chunks, the database system avoids expensive "vacuuming" operations when removing data, as the system can perform such operations by simply dropping chunks (internal tables and/or files), rather than deleting individual rows. For example, this removal may be the result of data deletions (e.g., based on automated data retention policies and procedures), or it may be the result of a large batch insert that fails or is interrupted (which the non-committed rows needing to subsequently be removed). At the same time, avoiding too-small chunks improves query performance by not needing to read additional tables and indexes from disk, or to perform query planning over a larger number of chunks.

The database system considers a few factors for determining a chunk's size. First, the database system maintains metadata that specify the number of partitions into which an additional partitioning field splits a particular time interval. For example, 10 machines each with 2 disks might use 20 partitions (or multiple partitions per server and/or disk). This implies that the key space of a particular field (e.g., a device ID, IP address, or location name) is divided into 20 ranges or sets. The database system then determines to which range (or partition) a particular value is associated by performing a lookup or comparison process. In one embodiment, the field is a string or binary value, and the database system splits the key space by prefix of the values of the field, then maps a value to one of these partitions based on the partition that shares the longest common prefix. Alternatively, the database system uses certain forms of hashing, such that the hash output's space is divided again into a particular number of ranges or sets (e.g., contiguous ranges, sets defined by splitting the entire hash output space, sets defined by taking the hash output space "mod" the number of nodes, sets defined by consistent hashing, etc.). The database system applies a hash function to the input value to yield an output value; the database system determines the range or set that includes the output value, which then corresponds to the partition to which the input value belongs. The database system may use a variety of functions in such a context, including hash functions (e.g., Murmur hashing, Pearson hashing, SHA, MD5, locality-sensitive hashing), the identity function (i.e., simply return the input), or some other prefixing or calculation on the input.

Second, once the number of partitions based on partitioning keys is determined—and in fact, this number can change over time due to elasticity, discussed below—then the time-duration of the chunk also determines its size. For a constant input rate and some given number of partitions, a chunk with a hour-long time interval will typically be much smaller than one with a day-long interval.

In one embodiment, the database system makes the time intervals static or manually configurable. Such an approach is appropriate if the data volumes to the system are relatively stable (and known), and this provides the database administrator or user with control over the database system's operation. But, such fixed time intervals may not work as

well as data volumes change—e.g., a time interval appropriate for a service pulling data from 100 devices is not appropriate when that system scales to 100,000 devices—or require care that the administrator or user change interval sizes over time (either to apply to future intervals or to split existing intervals into multiple chunks).

In one embodiment, the database system determines chunks' time intervals dynamically based on chunk sizes, rather than based on a fixed time interval. In particular, during insert time, the database system determines if a chunk is approaching or has exceeded some threshold size, at which time it "closes" the current chunk and creates a new chunk (e.g., by using the current time as the ending time of the current chunk and as the starting time of the new chunk).

This threshold size is given a default in software configuration, this default can be configured by the database system administrator, and this size can be changed by the administrator or the database system's logic during runtime (so that chunks in the same database system can have different threshold sizes). In an embodiment, the database system chooses the size as a function of the system's resources, e.g., based on the memory capacity of the server(s), which may also take into account the table schema to determine the amount of indexing that would be needed and its size requirements. This tuning takes into account realized or potential changes in the schema over time. For example, if indexes are added to many fields (columns), the amount of memory needed to store these fields changes, which leads the database system to use smaller chunks; if many fields are not indexed, the database system may account for these differently than a schema without any unindexed fields (as indexes may later be added to these fields to enable more efficient queries). Alternatively, recognizing that the database ultimately stores tables in files in the underlying file system that have a maximum size (e.g., 1 GB), the system ensures that the chunk size is smaller than this maximum size. In an embodiment, the size is chosen as a measured or estimated result of read/write performance on the chunk size.

In some embodiments, the database system creates a new chunk even when the current chunk size is less than some threshold (i.e., it is "approaching" the threshold, and has not yet exceeded or equaled it), in order to leave some "free space" for the possibility of out-of-time-order data that the database system must backfill into an older chunk. When writing to an older or "closed" chunk, different embodiments of the database system allow the chunk to grow arbitrarily large, create a new overlapping chunk just for the newly written excess data, or split the existing chunk into two, among other approaches. If overlapping chunks are created, the database system follows its policies for writing and reading to overlapping chunks.

In another embodiment, the database system determines a chunks' time intervals dynamically based on historical intervals and their sizes. In this case, new chunks are created with an end time, but that end time is automatically set by the database system based on the resulting size of earlier chunks that had a certain interval duration. For example, if the database system (or user or administrator) desires chunks of size approximation 1 GB, and the previous 12 hour chunk resulted in a chunk of size 1.5 GB, then the database might create a subsequent chunk of size 6 hours. The database system can continue to adapt the intervals of chunks during its operation, e.g., to account for changing data volumes per interval, to account for different target sizes, etc.

In some embodiments, the database determines chunks based on a hybrid of time intervals and threshold sizes. For

example, the database system (or administrator) specifies that a chunk have a pre-determined time interval—so that, as described above, the start and end time of a chunk are specified at creation time—but also that a chunk also have a maximum size in case the insert rate for that interval exceeds some amount. This approach avoids a problem with chunking based purely on fixed time-intervals in scenarios where system load per interval changes over time. If the chunk's size approaches or exceeds its maximum permitted threshold during the middle of the current time interval, the database system creates a new chunk that overlaps the same interval, or the database system switches to the use of a different time interval. For the former, both chunks represent the same interval, so inserts could choose to write to one of them (while reads query both of them). For the latter, the database system may change a chunk's time interval to something smaller, and create a new non-overlapping chunk to succeed it in time. As described earlier, such chunk management may be performed synchronously or asynchronously, e.g., a background task splits an over-large chunk into two chunks.

Such chunking may also limit the pre-determined time intervals to regular boundaries (e.g., 1 hour, 6 hours, 12 hours, 24 hours, 7 days, 14 days), rather than arbitrary ones (e.g., 11 minutes, 57 minutes). This embodiment causes chunk intervals to align well with periods of time on which data might be queried or deletions might be made, e.g., according to a data retention policy such as “delete data more than 12 hours old”. That way, the database system implements such policies by dropping entire chunks once their records are all at least 12 hours old, rather than partially deleting individual rows within chunks: dropping entire chunks (database tables) is much more efficient than deleting an equivalent number of rows within a table.

The database system selects these boundaries in a manner that the boundaries compose well, e.g., they are multiples of one another or are aligned in some other ways. The switching between various interval sizes is performed automatically by the database runtime (e.g., in response to changing data rates) or through configuration by a user or administrator. Similarly, rather than always closing a chunk and creating a new one based on an automated policy, an administrator may signal the database system to create a new chunk or chunk interval via a configuration command.

In one embodiment, the database system also applies such adaptation of the chunk's configuration to non-time dimension attributes that are used to define a chunk's ranges. For example, if a hypertable's partitioning is also performed on a field representing a device id, the database system can increase the number of partitions (sets of values) defined on this field from 10 to 20. Such a change, which may be performed automatically by the database system or through configuration by a user or administrator, can be used to increase hypertable performance. For example, if queries typically specify a single device id from which to SELECT data, the query's latency can be improved if the chunks that contain the specified device include information about a few other devices, which can be made to occur by increase the number of partitions over the device id field.

In another embodiment, the database system can employ different time intervals across different partitions. For example, if a hypertable's partitioning is also performed on a field representing a customer id (e.g., where each distinct customer id is a separate partition), then the database system may independently maintain different time intervals (when partitioning on the time attribute) for different customer ids. Such an approach can be beneficial if different customers

have very different insert and select query patterns, as well as different data retention needs.

In general, the database system employs a variety of methods for chunk management, given that there are multiple different goals and engineering trade-offs between approaches. These goals include optimizing sizes, aligning time intervals for dropping chunks while retaining data integrity, minimizing locking or other performance penalties due to mutability, avoiding arbitrary-sized intervals, creating chunk boundaries that are most advantageous for constraint exclusion, increasing system parallelism, improving query performance, and simplifying code, operation, and management complexity, among others. Different deployments of the database system may choose to use different approaches based on their setting and needs.

Adjusting Partitioning Policies in View of System Reconfiguration

The amount of data stored in a database systems increases over time. For example, large amount of time series data may be received by a database system and stored in database tables. Database systems often reconfigure the system to increase the storage capacity, for example, by adding storage devices. Conventional systems adapt to the change in the system configuration by moving data. For example, a system may get reconfigured as a result of addition of new servers and may move some chunks of data from existing servers to the new servers, in order to ensure that the new servers are bringing additional capacity to the system. As a result, a large amount of data is moved, thereby making the system reconfiguration an expensive and time-consuming process. This new configuration of participating servers is also referred to as a “view” which represents the set of servers and their configuration, such as the servers' capacity or number of disks. The ability of a system to adapt to changes in computing resources so as to be able to effectively use all available resources if referred to as elasticity.

Embodiments of the database system adapt to reconfiguration of the system without performing such data movement. In particular, the database system provides elasticity by creating a new set of chunks and partitioning when the database system is reconfigured for increasing the storage capacity. The database system may use a different partitioning policy for the new set of chunks that are created after the system is reconfigured. For example, if the previous partitioning policy created 20 partitions for 10 servers, the new partitioning policy might create 30 partitions to take into account 5 new servers that are added to the database system. In another example, the previous partitioning policy may create 20 partitions to place 5 partitions on each of 4 servers, but when an additional 1 server is added, the new partitioning policy may then place 4 partitions on each of the 5 servers. In some embodiments, the database system distributes a plurality of chunks created such that new servers are assigned more chunks from the plurality of chunks than existing servers. This allows better balancing of load across the servers. In another embodiment, new servers are assigned larger chunks compared to chunks assigned to existing servers. Larger chunks have configuration that allows them to potentially store more data than smaller chunks. Data can still be read or written to previously created chunks or the newly created chunks. Because writes to time-series datasets are typically made to the latest time interval, and many query workloads also process recent data, load balancing across the new set of servers is still maintained, even without moving the existing data.

FIGS. 7(A-B) illustrate partitioning of data of a database table to adapt to the addition of locations to the database system according to an embodiment of the invention.

As illustrated in FIG. 7(A), the database system **110** can have a plurality of storage locations **710a**, **710b**, **710c**, and **710d**. FIG. 7 illustrates the distribution of data of a database table with attributes comprising a time attribute and a space attribute (recall that we use the term “space” partitioning to signify any partitioning over a non-time attribute). In response to requests to insert records in the database table, the database system **110** distributes data of the database table according to a partitioning policy that assigns chunks **210** to locations **710**. In the example, configuration shown in FIG. 7(A), the database system **110** creates a plurality of chunks including **210a**, **210b**, **210c**, and **210d** and assigns one chunk to each location. The chunks are distributed across the locations of the database system **110** along the time and space attributes. Accordingly, each chunk is associated with a time range and a space range and stores records that have time and space attributes that lie within the time and space ranges of the chunk. In the example configuration shown in FIG. 7, each of the chunks **210a**, **210b**, **210c**, and **210d** is associated with the same range of time attribute, i.e., [0, 6] but has a different range of the space attribute. For example, chunk **210a** has space range [A, F], the chunk **210b** has space range [G, L], the chunk **210c** has space range [M, S], and the chunk **210d** has space range [T, Z].

FIG. 7(B) shows the partitioning of the chunks of the database table after some time has passed, such that the database system has received requests to insert records with a time attribute later than 6. In response to receiving requests to insert records with a time attribute of 7, for example, the database system creates a new plurality of chunks, **201e**, **210f**, **210g**, and **210h**. The new plurality of chunks are distributed across the locations according to the same partitioning policy as above. According to this partitioning policy, each chunk from the new plurality of chunks is associated with a new time range [7, 15]. In this illustration, the chunks stored in the same location have the same space range. For example, both chunks **210a** and **210e** assigned to location **710a** have the space range [A, F], both chunks **210b** and **210f** assigned to location **710b** have the space range [G, L], and so on. The database system could also assign chunks with different time intervals but the same space range to different locations.

FIG. 7(C) shows the partitioning of the chunks of the database table after a new location **710e** is added to the database system **110**. As a result, the database system **110** has a plurality of locations that include locations **710a**, **710b**, **710c**, **710d**, and **710e**. Although FIG. 7 shows a single location being added to the database system, more than one locations may be added to increase the storage capacity of the database system **110**. In response to addition of the new location, the database system **110** uses a new partitioning policy to distribute records across the locations. Accordingly, in response to receiving subsequent insert requests, e.g., with values for dimension attributes that do not map to any of the existing chunks, the database system **110** creates a plurality of chunks including **210i**, **210j**, **210k**, **210l**, **210m**, and **210n**. The chunks **210i**, **210j**, **210k**, **210l** are mapped to the locations **710a**, **710b**, **710c**, **710d**, and chunks **210m** and **210n** are both mapped to the new location **710e**. In other embodiments, the database system may assign more or fewer chunks to the new locations that are added. Accordingly, subsequent records received are distributed according to the new partitioning policy. In the embodiment illustrated in FIG. 7, the database system **110** does not move any data

that was stored in the chunks that were created before the new locations were added. However, the chunks that are created responsive to the addition of the new locations are distributed according to a new partitioning policy that balances storage of data across all available locations. In the example, shown in FIG. 7(C), more chunks are assigned to the new location(s) since the storage and computing resources of the new locations are likely to be underutilized compared to the existing locations that have previously stored data. However, over time, as additional data gets stored on the new locations, the utilization gap between the new locations and existing locations reduces without having to move any data from the existing locations to the new locations.

As illustrated in FIG. 7(C), the new partitioning policy creates a plurality of chunks that has more chunks after new locations are added. Accordingly, each space range is smaller in the new partitioning policy compared to the space ranges of the partitioning policy used before addition of the new locations.

In another embodiment, the database system assigns a larger fraction of new data to the new locations not by assigning a larger number of chunks to those locations, as shown in FIG. 7(C), but by assigning chunks with dimension ranges that have a larger set of values. For example, rather than having chunk **210m** with the space range [Q, U] and chunk **210n** with the space range [V, Z], the database system could create a single chunk assigned to location **710e** with a space range [Q, Z].

In some embodiments, when the database system **110** detects that new locations are being added to the database system, the database system **110** dynamically changes the partitioning based on the new storage configuration. In other embodiments, the partitioning policy is configured by a user, for example, a database system administrator.

A partitioning policy determines how new chunks are created and assigned to locations for storing them. For example, if a partitioning policy is being enforced and new chunks need to be created (for example, to insert records that cannot be inserted in existing chunks), a plurality of chunks may be created and distributed according to the partitioning policy. The partitioning policy may specify various aspects of creation of new chunks including the number of chunks being created, the configurations of individual chunks being created (the configuration comprising the sets of values of different dimension attributes for each chunk), and the mapping of the chunks to locations.

The partitioning policy may store information specifying various aspects of the chunk creation/distribution as metadata, for example, the mapping from chunks to locations may be stored using a mapping table that explicitly stores locations for each chunk being created. Alternatively, the partitioning policy may specify various aspects of chunk creation/distribution using instructions, for example, the partitioning policy may specify mapping from chunks to locations using a function (or a set of instructions) that determines a location for a chunk given the chunk configuration and potentially other system information as input. Different partitioning policies may specify different mapping functions (or sets of instructions). Alternatively, different partitioning policies may use the same mapping function (or sets of instructions) but pass different parameter values as input. Such mapping functions (or sets of instructions) may include random selection, round-robin selection, hash-based selection, selection based on the number, size, or age of chunks being stored, selection based on the age of when the location was added to the database system, load balanc-

ing strategies based on server resources (including insert or query rates, CPU capacity, CPU utilization, memory capacity, free memory, etc.), load balancing strategies based on disk resources (including total disk capacity, unused disk space disk, disk IOPS capacity, disk IOPS use, etc.), and other criteria or algorithmic approaches, as well as some combination thereof. A partitioning policy may use a combination of the above techniques.

In an embodiment, a partitioning policy specifies the size of the plurality of chunks being created. The size of the plurality of chunks may represent the number of chunks in the plurality of chunks being created. Alternatively, the size of the plurality of chunks may represent the aggregate size of chunks in the plurality of chunks being created, where the size of each chunk represents a measure of the amount of data that can potentially be stored in the chunk. The size of a chunk is determined based on the configuration of the chunk comprising the sets of values of the different dimension attributes for records stored in the chunk. For example, the database system may create larger or smaller chunks by specifying larger/smaller ranges (or sets of values) for dimension attributes respectively.

In some embodiments, the database system **110** moves existing data under certain scenarios. For example, the database system may enforce a policy that aligns chunks to specific time intervals. Accordingly, the creation of new chunks at a time based on the time that new locations are added may result in violation of such policy. For example, the database system may enforce a standard that chunks have a time range of 12 hours. However, if the addition of new locations to the database system occurs at 3 hours into a 12-hour time interval, the database system would either not be able to incorporate the new locations for another 9 hours, or would have to maintain some chunks with 3 hours intervals. Thus, in certain scenarios, for example, if the amount of data stored in each chunk that is currently being populated is below a threshold amount, the database system moves or reallocates existing chunks rather than create new ones responsive to addition of new location. Accordingly, the database system moves data of the set of chunks being currently populated with records across a new set of chunks distributed across the new plurality of locations and continues adding records to the new set of chunks.

In another embodiment, the database system delays enforcement of the new partitioning policy based on the new locations added until the time matches well with chunk alignments. This delayed action can be used both when adding new servers, removing servers in a planned manner, or even on server crashes (if the system already replicates chunks between multiple servers for high availability). For example, if the system already has chunks with time ranges that extend until midnight, and the reconfiguration time is at 11 pm, the database system may not create chunks based on the new partitioning policy for 1 hour (e.g., until a record is inserted with a time attribute after midnight), but the reconfiguration will have an effect when a new set of chunks is created. In such a scenario, the existing chunks are not reconfigured and only the new chunks are allocated over the new set of servers. However, the time range of the chunks is the same before and after the addition of the new locations.

FIG. 8 shows a flowchart illustrating the process of modifying a data partitioning policy of a database system in response to the addition of new locations to the database system, in accordance with an embodiment. The database system **110** includes a plurality of locations, referred to as the first plurality of locations. The database system **110** receives **810** requests to insert records in a hypertable. The

database system distributes the chunks in accordance with a first partitioning policy **P1**. Accordingly, the database system **110** creates **820** a plurality of chunks and distributes them across the first plurality of locations. For example, if the database system has 5 locations, the database system **110** may create 20 chunks and store 4 chunks in each location. The database system **110** distributes the chunks based on dimension attributes of the records including at least a time attribute. The partitioning policy specifies various aspects of chunk/creation and distribution including the number of chunks that may be created, the configuration of the chunks, and the mapping of the chunks to locations. The database system may repeat the steps **810** and **820** multiple times, for example, until the database system **110** is reconfigured to change the number of locations.

The database system **110** receives an indication of the addition of one or more new locations. For example, a new location may be a storage device that is added by a system administrator to an existing server of the database system. Alternatively, a new location may be a new server comprising one or more storage devices that is added to the database system for storing as well as processing data. As another example, a location may be storage device of a remote system on which the database system **110** is allowed to store data, for example, a cloud-based storage device. The indication of the addition of one or more new locations that the database system receives may identify a specific storage device that is added to the database system or may identify a server that is added to the database system.

In an embodiment, the database system **110** receives the indication of addition of a location by performing a check of all peripheral devices and servers that can be reached by one or more database system nodes **310**. In other embodiments, the database system **110** receives the indication by receiving a message from a new location, by a command executed by a database user or administrator. The addition of the locations to the database system causes the database system **110** to have a second plurality of locations that is more than the number of locations in the first plurality of locations. The indication of addition of the one or more locations is associated with a reconfiguration time, for example, the time that the indication is received or the time when the addition of the one or more new locations was completed.

Subsequent to receiving the indication of the addition of one or more new locations, the database system receives insert requests. The database system **110** creates **840** a second plurality of chunks, for example, if the records in the insert requests received cannot be inserted in existing chunks. The database system **110** creates the second plurality of chunks and assigns them to locations based on a second partitioning policy **P2**. The second partitioning policy **P2** maps the second plurality of chunks to the second plurality of locations, for example, as illustrated in FIG. 7(C). The chunks may be uniformly distributed across the second plurality of locations. Alternatively, the number or partition ranges of chunks assigned to the new locations may be greater than the number or partition ranges of chunks assigned to the existing locations. For example, more chunks from the second plurality of chunks may be assigned to the new locations compared to the existing locations. Alternatively, chunks configured to store more data may be assigned to new locations compared to the existing locations. A chunk **C1** may be configured to store more data compared to a chunk **C2** by specifying for chunk **C1**, a set of values for a dimension attribute that has more elements compared to the set of values for the same dimension attribute for chunk **C2**.

For example, the time attribute for chunk C1 may be specified to have a larger time range compared to the time attribute for chunk C2.

The database system 110 subsequently receives 850 requests to insert data in the database table. The database system 110 stores 860 the received records into chunks based on the dimension attributes of the records. The records may be inserted in chunks created either based on the first partitioning policy or the second partitioning policy as further described herein in connection with FIGS. 9-12. The database system 110 identifies a reconfiguration time T associated with the addition of the new locations to the database system.

In an embodiment, the database system inserts records into chunks based on a time attribute of the record. Accordingly, even though a new partitioning policy is defined, the database system may receive insert requests and create chunks based on a previous partitioning policy. For example, the database system may receive some records very late (i.e., the time they are received may be significantly after the values of the records' time attribute), for example, due to delay caused by network or other resources. The database system may create chunks based on an older partitioning policy for storing these records. Accordingly, the database system may enforce multiple partitioning policies at the same time, depending on the data of the records that are received and need to be inserted in a hypertable.

FIG. 9 illustrates selection of partitioning policy for creating chunks based on time attribute of the record, according to an embodiment. Accordingly, independent of the time that the insert request is received, if insert requests are received with records having a time attribute value that is before the reconfiguration time T, any new chunks created for storing the records are created based on the first partitioning policy. FIG. 9 shows a timeline 900 and various events along the time line. For example, the database system initially has three locations (disks) 910a, 910b, and 910c and creates chunks according to partitioning policy P1. At reconfiguration time T, a new location 910d is added to the database system 110. However, if insert requests received after reconfiguration time T have time attribute values that are before reconfiguration time T, the database system creates chunks for storing the records (if none of the existing chunks can store the records) according to the first partitioning policy P1. Furthermore, if insert requests received after reconfiguration time T have time attribute values that are after the reconfiguration time T, the database system creates chunks for storing the records (if none of the existing chunks can store the records) according to the second partitioning policy P2. Accordingly, the time interval T1 during which chunks are created according to the first partitioning policy P1 can extend after the reconfiguration time T. Time interval T2 indicates the time during which chunks are created according to the second partitioning policy P2.

FIG. 10 shows a flowchart of the process for selection of partitioning policy for creating chunks based on time attribute of the record, according to an embodiment. The database system invokes the procedure shown in FIG. 10 if the database system determines for a record being inserted that the record cannot be stored in any existing chunk and a new chunk needs to be created. The database system 110 determines 1010 the value of the time attribute of a record received for inserting in the database table. The database system 110 compares 1020 the value of the time attribute of the record with the reconfiguration time T. If the database system 110 determines that the time attribute of the record

is less than the reconfiguration time T, the database system 110 creates a chunk 1030 based on the first partitioning policy P1. If the database system 110 determines that the time attribute of the record is greater than (or equal to) the reconfiguration time T, the database system 110 creates 1040 a chunk based on the second partitioning policy P2. The record is stored 1050 in the chunk that is created.

FIG. 11 illustrates selection of partitioning policy for creating chunks based on time of receipt of a record by the database system, according to an embodiment. FIG. 11 shows a timeline 1100 and various events along the time line. For example, the database system initially has three locations (disks) 1110a, 1110b, and 1110c and creates chunks according to partitioning policy P1. At reconfiguration time T, a new location 1110d is added to the database system 110. The database system selects the partitioning policy for creating chunks based on the time of arrival of the insert request (assuming no existing chunks can be used for storing records that are received for inserting in the hypertable). Accordingly, after reconfiguration time T (i.e., during time interval T2), chunks are created according to the second partitioning policy P2 whereas before reconfiguration time T (i.e., during time interval T1), chunks are created according to the first partitioning policy P1. Accordingly, the partitioning policy selected for creating chunks is selected independently of the value of the time attribute of the records being inserted. For example, if for any reason records having time attribute values that correspond to time occurring before reconfiguration time T arrive late, i.e., after reconfiguration time T, the database system creates chunks according to the second partitioning policy P2 for storing the records. Accordingly, records with time attribute value less than reconfiguration time T can be stored in chunks created according to either partitioning policy P1 or P2.

In some embodiments, the database system continues to insert records into a chunk that was created before reconfiguration time T even if the insert request arrives after reconfiguration time T so long as the time attribute of the record corresponds to the time range for the chunk. In other embodiments, the database system modifies an existing chunk that was created according to the first partitioning policy P1 so as to reduce the time range (if necessary) to correspond to the latest record inserted into the chunk. For example, if the insert request's arrival time is 5:30 am and the chunk's current time range is until noon, the database system identifies the record with the highest value for its time attribute in that chunk. Assuming that the record with the highest time value in that chunk has a time of 5:45 am, the database system modifies the end of the chunk's time range to a time greater than or equal to 5:45 am, for example, 6 am. Subsequently, if the database system receives a record at time greater than 6 am, the database system creates a new chunk according to the new partitioning policy P2 starting at 6 am.

In some embodiments, the database system may create overlapping chunks as a result of reconfiguration of the system. The database system enforces a policy that after reconfiguration of the system, the database system does not insert records in chunks created based on the first partitioning policy P1. As a result, after reconfiguration of the system, the database system creates a new chunk for storing a record based on partitioning policy P2, even if there is an existing chunk created based on policy P1 that maps to the dimension attributes of the record. As a result, a record having a particular dimension attribute could potentially be stored in a chunk C1 created based on the first partitioning policy P1 or in a chunk C2 created based on the second

partitioning policy P2. As a result, chunks C1 and C2 are overlapping such that a record could map to both chunks C1 and C2. If the database system subsequently receives queries that process a particular record R, the database system 110 determines whether the record R was stored in a chunk created based on the first partitioning policy P1 or the second partitioning policy P2. Accordingly, the database system 110 may have to check two possible chunks to determine where the record R is stored.

In some embodiments, the database system 110 creates the new chunks that overlap old chunks in terms of the time range used for partitioning the records. As a result, even after creation of a new set of chunks responsive to the addition of new locations, the database system may insert records into old chunks that were created before the addition of the locations. While this may involve the old chunks (from the old view) continuing to see some fraction of new inserts—although this can be mitigated based on the insert policy for overlapping chunks, e.g., one such policy prefers inserting new records to the smaller-sized chunk—this overlap will not continue into future intervals. For example, continuing with the above example, when the database system creates the new chunks 9 hours into the existing chunks' interval, it sets the start and end times for the new chunks to be the same as the existing chunks (i.e., 9 hours ago and 3 hours hence). But, because the database system can employ a policy to write to smaller-sized chunks, for example, inserts will be made to the new chunks rather than the existing ones, even though the two sets have overlapping time periods.

In embodiments of the database system that use a purely size-based approach to determining when to close a chunk, these time interval issues do not arise, and the database system then simply closes the existing chunks (even when their size at the time of system reconfiguration may be smaller than the standard threshold size) and creates new ones using the new partitioning policy.

Because the new view may maintain a different set of partitions, the database system may maintain additional metadata that associates each of these reconfigurations into an "epoch." In particular, each epoch may be associated with various information, including a time period, the set of partitions, and a system view. Then, as described above, in order to determine a hypertable's partitions at a particular time, the database system might need to first determine the epoch associated with the time, then determine the partitions associated with this epoch. This process is described above in the context of an insert method that the database system employs.

#### Architecture of Computer for a Database System

FIG. 12 is a high-level block diagram illustrating an example of a computer 1200 for use as one or more of the entities illustrated in FIG. 1, according to one embodiment. Illustrated are at least one processor 1202 coupled to a memory controller hub 1220, which is also coupled to an input/output (I/O) controller hub 1222. A memory 1206 and a graphics adapter 1212 are coupled to the memory controller hub 1222, and a display device 1218 is coupled to the graphics adapter 1212. A storage device 1208, keyboard 1210, pointing device 1214, and network adapter 1216 are coupled to the I/O controller hub. The storage device may represent a network-attached disk, local and remote RAID, or a SAN (storage area network). A storage device 1208, keyboard 1210, pointing device 1214, and network adapter 1216 are coupled to the I/O controller hub 1222. Other embodiments of the computer 1200 have different architectures. For example, the memory is directly coupled to the processor in some embodiments, and there are multiple

different levels of memory coupled to different components in other embodiments. Some embodiments also include multiple processors that are coupled to each other or via a memory controller hub.

The storage device 1208 includes one or more non-transitory computer-readable storage media such as one or more hard drives, compact disk read-only memory (CD-ROM), DVD, or one or more solid-state memory devices. The memory holds instructions and data used by the processor 1202. The pointing device 1214 is used in combination with the keyboard to input data into the computer 1200. The graphics adapter 1212 displays images and other information on the display device 1218. In some embodiments, the display device includes a touch screen capability for receiving user input and selections. One or more network adapters 1216 couple the computer 1200 to a network. Some embodiments of the computer have different and/or other components than those shown in FIG. 12. For example, the database system can be comprised of one or more servers that lack a display device, keyboard, pointing device, and other components, while a client device acting as a requester can be a server, a workstation, a notebook or desktop computer, a tablet computer, an embedded device, or a handheld device or mobile phone, or another type of computing device. The requester to the database system also can be another process or program on the same computer on which the database system operates.

The computer 1200 is adapted to execute computer program modules for providing functionality described herein. As used herein, the term "module" refers to computer program instructions and/or other logic used to provide the specified functionality. Thus, a module can be implemented in hardware, firmware, and/or software. In one embodiment, program modules formed of executable computer program instructions are stored on the storage device, loaded into the memory, and executed by the processor.

#### Additional Considerations

In time-series workloads, writes are typically made to recent time intervals, rather than distributed across many old ones. This allows the database system 110 to efficiently write batch inserts to a small number of tables as opposed to performing many small writes across one giant table. Further, the database systems' clustered architecture also takes advantage of time-series workloads to recent time intervals, in order to parallelize writes across many servers and/or disks to further support high data ingest rates. These approaches improve performance when employed on various storage technologies, including in-memory storage, hard drives (HDDs), or solid-state drives (SSDs).

Because chunks are right-sized to servers, and thus the database system does not build massive single tables, the database system avoids or reduces swapping its indexes to disks for recent time intervals (where most writes typically occur). This occurs because the database system maintains indexes local to each chunk; when inserting new records into a chunk, only that chunks' (smaller) indexes need to be updated, rather than a giant index built across all the hypertable's data. Thus, for chunks associated with recent time intervals that are regularly accessed, particularly if the chunks are sized purposefully, the chunks' indexes can be maintained in memory. Yet the database system can still efficiently support many different types of indexes on different types of columns (e.g., based on what is supported by each node's database engine, such as PostgreSQL), including B-tree, B+-tree, GIN, GiST, SP-GiST, BRIN, Hash, LSM Tree, fractal trees, and other types of indexes.

The database system combines the transparent partitioning of its hypertable abstraction with a number of query optimizations. These optimizations include those which serve to minimize the number and set of chunks that must be contacted to satisfy a query, to reduce the amount of records that are transferred back from a query that touches a chunk, to specify whether raw records or aggregates results are transferred back from a chunk, and so forth.

Common queries to time-series data include (i) slicing across time for a given object (e.g., device id), slicing across many objects for a given time interval, or (iii) querying the last reported data records across (a subset of) all objects or some other distinct object label. While users perform these queries as if interacting with a single hypertable, the database system leverages internally-managed metadata to only query those chunks that may possibly satisfy the query predicate. By aggressively pruning many chunks and servers to contact in its query plan—or during execution, when the system may have additional information—the database system improves both query latency and throughput.

Similarly, for items like unique devices, users, or locations, the database system may receive queries like “select the last K readings for every device.” While this query can be natively expressed in SQL using a “SELECT DISTINCT” query (for finding the first or last single value per distinct item) or via windowing functions (for finding K such values), such a query can turn into a full table scan in many relational databases. In fact, this full table scan could continue back to the beginning of time to capture “for every device”, or otherwise either sacrifice completeness with some arbitrarily-specified time range or involve a large WHERE clause or JOIN against some set of devices of interest (which may be maintained in a manual or automated fashion).

In some embodiments, the database system maintains additional metadata about a hypertable’s fields in order to optimize such queries. For example, the database system records information about every distinct (different) value for that field in the database (e.g., the latest row, chunk, or time interval to which it belongs). The database system uses this metadata along with its other optimizations, so that such queries for distinct items avoid touching unnecessary chunks, and perform efficiently-indexed queries on each individual chunk. The decision to maintain such metadata might be made manually or via automated means for a variety of reasons, including based on a field’s type, the cardinality of the field’s distinct items, query and workload patterns, and so forth.

The database system may perform other query optimizations that benefit both single-node and clustered deployments. When joining data from multiple tables (either locally or across the network, e.g., via foreign data wrappers), traditional databases may first select all data matching the query predicate, optionally ORDER the data, then perform the requested LIMIT. Instead, the database system **110** first performs the query and post-processing (e.g., ORDER and LIMIT) on each chunk, and only then merges the resulting set from each chunk (after which it performs a final ordering and limit).

The database system **110** uses LIMIT pushdown for non-aggregate queries to minimize copying data across the network or reading unnecessary data from tables. The database system also pushes down aggregations for many common functions (e.g., SUM, AVG, MIN, MAX, COUNT) to the servers on which the chunks reside. Primarily a benefit for clustered deployments, this distributed query optimization greatly minimizes network transfers by performing

large rollups or GROUP\_BYs in situ on the chunks’ servers, so that only the computed results need to be joined towards the end of the query, rather than raw data from each chunk. In particular, each node in the database system performs its own partial aggregation, and then only return that result to the requesting node.

For example, if the query to the database system requests some MAX (maximum value), then the first node processing the hypertable query sends MAX queries to other nodes; each receiving node performs the MAX on its own local chunks before sending the result back to the first node. This first node computes the MAX of these local maximum values, and returns this result. Similarly, if the hypertable query asks for the AVG (average), then the first node sends queries to other servers that ask for the sum and count of some set of rows. These nodes can return their sums and counts to the first node, which then computes the total average from these values (by dividing the sum of sums by the sum of counts).

The database system computes joins between hypertables and standard relational tables. These standard tables can be stored either directly in the database system or accessed from external databases, e.g., via foreign data wrappers.

The database system **110** performs joins between two hypertables, including in a number of ways involving distributed optimizations, e.g., distributed joins. Such optimizations include those using hash-based partitioning, as well as those that carefully minimize data copying by only sending data from one hypertable’s chunks to the servers with the other’s chunks according to the join being performed, optionally leveraging the metadata associated with the chunk. Such optimizations also include placing the chunks of hypertables that will be regularly joined on servers in a way that like keys or key ranges are commonly collocated on the same server, to minimize sending data over the network during joins.

The database system allows for easily defining data retention policies based on time. For example, administrators or users can use explicit commands or configure the system to cleanup/erase data more than X weeks old. The system’s chunking also helps make such retention policies more efficient, as the database system then just drops entire chunks (internal data tables) that are expired, as opposed to needing to delete individual rows and aggressively vacuum the resulting tables, although the database system does support such row-based deletions.

For efficiency, the database system enforces such data retention policies lazily. That is, individual records that are older than the expiry period might not be immediately deleted, depending upon policy or configuration. Rather, when all data in a chunk becomes expired, then the entire chunk is dropped. Alternatively, the database system uses a hybrid of dropping chunks and deleting individual rows when performing data deletions or adhering to data retention policies.

The foregoing description of the embodiments of the invention has been presented for the purpose of illustration; it is not intended to be exhaustive or to limit the invention to the precise forms disclosed. Persons skilled in the relevant art can appreciate that many modifications and variations are possible in light of the above disclosure.

Some portions of this description describe the embodiments of the invention in terms of algorithms and symbolic representations of operations on information. These algorithmic descriptions and representations are commonly used by those skilled in the data processing arts to convey the substance of their work effectively to others skilled in the



art. These operations, while described functionally, computationally, or logically, are understood to be implemented by computer programs or equivalent electrical circuits, microcode, or the like. Furthermore, it has also proven convenient at times, to refer to these arrangements of operations as modules, without loss of generality. The described operations and their associated modules may be embodied in software, firmware, hardware, or any combinations thereof.

Any of the steps, operations, or processes described herein may be performed or implemented with one or more hardware or software modules, alone or in combination with other devices. In one embodiment, a software module is implemented with a computer program product comprising a computer-readable medium containing computer program code, which can be executed by a computer processor for performing any or all of the steps, operations, or processes described.

Embodiments of the invention may also relate to an apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, and/or it may comprise a general-purpose computing device selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a tangible computer readable storage medium or any type of media suitable for storing electronic instructions, and coupled to a computer system bus. Furthermore, any computing systems referred to in the specification may include a single processor or may be architectures employing multiple processor designs for increased computing capability.

Finally, the language used in the specification has been principally selected for readability and instructional purposes, and it may not have been selected to delineate or circumscribe the inventive subject matter. It is therefore intended that the scope of the invention be limited not by this detailed description, but rather by any claims that issue on an application based hereon. Accordingly, the disclosure of the embodiments of the invention is intended to be illustrative, but not limiting, of the scope of the invention.

What is claimed is:

1. A computer-implemented method comprising:

receiving, by a database system, one or more insert requests for inserting rows in a hypertable, the one or more insert requests specifying a set of input records for inserting in the hypertable, each record from the set of records having a plurality of attributes including a set of dimension attributes, the set of dimension attributes including a time attribute, wherein the hypertable represents a database table partitioned into a plurality of chunks along the set of dimension attributes, each chunk having a configuration associated with a set of values corresponding to each dimension attribute, such that, for each record stored in the chunk, and for each dimension attribute of the record, the value of the dimension attribute of the record maps to a value from the set of values for that dimension attribute as specified by the chunk;

for at least an input record from the set of input records, storing the input record in a new chunk, wherein the new chunk is created subsequent to receiving at least one of the one or more insert requests, wherein the input record is stored in the new chunk based on a determination that the dimension attributes of the input record match the configuration of the new chunk, wherein the dimension attributes of the input record do not match the configurations of the chunks of the hypertable created before the new chunk;

storing one or more policies specifying actions associated with chunks matching criteria related to configurations of the chunks;

selecting one or more chunks subsequent to storing the input record, the one or more chunks including the new chunk, the selecting according to criteria specified in a policy; and

performing an action associated with the one or more selected chunks, the action specified by the policy.

2. The computer-implemented method of claim 1, wherein the selected one or more chunks are distributed across two or more locations, and wherein each location corresponds to one of:

a storage device added to a server from the set of servers used by the database system;

a storage device belonging to a new server, wherein the new server is added to the set of servers used by the database system; or

a network attached storage device of a remote server made accessible to a server from the set of servers used by the database system.

3. The computer-implemented method of claim 1, wherein selecting one or more chunks according to criteria specified in the policy configuration comprises, for each of the one or more chunks, comparing the time attribute of the chunk with a threshold value.

4. The computer-implemented method of claim 1, wherein the criteria specified in the policy configuration comprises age of the chunk, wherein age of a chunk is determined based on the set of values associated with the time attribute of the chunk.

5. The computer-implemented method of claim 1, wherein the action associated with a selected chunk comprises deleting the selected chunk.

6. The computer-implemented method of claim 1, wherein the action associated with a selected chunk comprises performing one or more of:

adding an index to the selected chunk,

dropping an index from the selected chunk,

reindexing the selected chunk, or

adding a constraint on the chunk.

7. The computer-implemented method of claim 1, wherein the action associated with a selected chunk comprises moving the selected chunk from a first location to a second location.

8. The computer-implemented method of claim 7, wherein the second location is determined based on the policy or the configuration of the selected chunk.

9. The computer-implemented method of claim 7, wherein the second location is determined based on one or more of: access time of location, storage capacity of location, cost of location, or performance of location.

10. The computer-implemented method of claim 1, wherein the action associated with a selected chunk comprises moving one or more records of the chunk to another chunk and updating the configurations of the selected chunk and the other chunk based on the movement of the one or more records.

11. The computer-implemented method of claim 1, wherein a selected chunk has a plurality of replicas, each of the plurality of replicas storing a copy of a particular record.

12. The computer-implemented method of claim 11, wherein the action is performed on one or more replicas from the plurality of replicas.

13. The computer-implemented method of claim 11, wherein each of the plurality of replicas is stored in a different location.

14. The computer-implemented method of claim 11, wherein the action comprises moving a replica from a first location to a second location, where the second location has different characteristics than the first location.

15. The computer-implemented method of claim 11, wherein the number of replicas of a chunk are determined based on the configuration of the chunk.

16. The computer-implemented method of claim 11, wherein the number of replicas of the chunk is determined based on a set of values associated with the time attribute for the chunk.

17. A non-transitory computer readable storage medium storing instructions that when executed by a computer processor, cause the computer processor to perform steps comprising:

receiving, by a database system, one or more insert requests for inserting rows in a hypertable, the one or more insert requests specifying a set of input records for inserting in the hypertable, each record from the set of records having a plurality of attributes including a set of dimension attributes, the set of dimension attributes including a time attribute, wherein the hypertable represents a database table partitioned into a plurality of chunks along the set of dimension attributes, each chunk having a configuration associated with a set of values corresponding to each dimension attribute, such that, for each record stored in the chunk, and for each dimension attribute of the record, the value of the dimension attribute of the record maps to a value from the set of values for that dimension attribute as specified by the chunk;

for at least an input record from the set of input records, storing the input record in a new chunk, wherein the new chunk is created subsequent to receiving at least one of the one or more insert requests, wherein the input record is stored in the new chunk based on a determination that the dimension attributes of the input record match the configuration of the new chunk, wherein the dimension attributes of the input record do not match the configurations of the chunks of the hypertable created before the new chunk;

storing one or more policies specifying actions associated with chunks matching criteria related to configurations of the chunks;

selecting one or more chunks subsequent to storing the input record, the one or more chunks including the new chunk, the selecting according to criteria specified in a policy; and

performing an action associated with the one or more selected chunks, the action specified by the policy.

18. The non-transitory computer readable storage medium of claim 17, wherein the selected one or more chunks are distributed across two or more locations, and wherein each location corresponds to one of:

a storage device added to a server from the set of servers used by the database system;

a storage device belonging to a new server, wherein the new server is added to the set of servers used by the database system; or

a network attached storage device of a remote server made accessible to a server from the set of servers used by the database system.

19. The non-transitory computer readable storage medium of claim 17, wherein instructions for selecting one or more chunks according to criteria specified in the policy configuration comprise instructions for:

for each of the one or more chunks, comparing the time attribute of the chunk with a threshold value.

20. The non-transitory computer readable storage medium of claim 17, wherein the criteria specified in the policy configuration comprises age of the chunk, wherein age of a chunk is determined based on the set of values associated with the time attribute of the chunk.

21. The non-transitory computer readable storage medium of claim 17, wherein the action associated with a selected chunk comprises deleting the selected chunk.

22. The non-transitory computer readable storage medium of claim 17, wherein the action associated with a selected chunk comprises performing one or more of:

adding an index to the selected chunk,  
dropping an index from the selected chunk,  
reindexing the selected chunk, or  
adding a constraint on the chunk.

23. The non-transitory computer readable storage medium of claim 17, wherein the action associated with a selected chunk comprises moving the selected chunk from a first location to a second location.

24. The non-transitory computer readable storage medium of claim 23, wherein the second location is determined based on the policy or the configuration of the selected chunk.

25. The non-transitory computer readable storage medium of claim 23, wherein the second location is determined based on one or more of: access time of location, storage capacity of location, cost of location, or performance of location.

26. The non-transitory computer readable storage medium of claim 17, wherein the action associated with a selected chunk comprises moving one or more records of the chunk to another chunk and updating the configurations of the selected chunk and the other chunk based on the movement of the one or more records.

27. The non-transitory computer readable storage medium of claim 17, wherein a selected chunk has a plurality of replicas, each of the plurality of replicas storing a copy of a particular record.

28. The non-transitory computer readable storage medium of claim 27, wherein the action is performed on one or more replicas from the plurality of replicas.

29. The non-transitory computer readable storage medium of claim 27, wherein each of the plurality of replicas is stored in a different location.

30. The non-transitory computer readable storage medium of claim 27, wherein the action comprises moving a replica from a first location to a second location, where the second location has different characteristics than the first location.

31. The non-transitory computer readable storage medium of claim 27, wherein the number of replicas of a chunk are determined based on the configuration of the chunk.

32. The non-transitory computer readable storage medium of claim 27, wherein the number of replicas of the chunk is determined based on a set of values associated with the time attribute for the chunk.

33. A computer system comprising:  
a computer processor; and

a non-transitory computer readable storage medium storing instructions that when executed by the computer processor, cause the computer processor to perform steps comprising:

receiving, by a database system, one or more insert requests for inserting rows in a hypertable, the one or more insert requests specifying a set of input records for inserting in the hypertable, each record from the set of records having a plurality of attributes including a set of dimension attributes, the set of dimension

39

attributes including a time attribute, wherein the hypertable represents a database table partitioned into a plurality of chunks along the set of dimension attributes, each chunk having a configuration associated with a set of values corresponding to each dimension attribute, such that, for each record stored in the chunk, and for each dimension attribute of the record, the value of the dimension attribute of the record maps to a value from the set of values for that dimension attribute as specified by the chunk;

for at least an input record from the set of input records, storing the input record in a new chunk, wherein the new chunk is created subsequent to receiving at least one of the one or more insert requests, wherein the input record is stored in the new chunk based on a determination that the dimension attributes of the input record match the configuration of the new chunk, wherein the dimension attributes of the input record do not match the configurations of the chunks of the hypertable created before the new chunk;

storing one or more policies specifying actions associated with chunks matching criteria related to configurations of the chunks;

selecting one or more chunks subsequent to storing the input record the one or more chunks including the new chunk, the selecting according to criteria specified in a policy; and

performing an action associated with the one or more selected chunks, the action specified by the policy.

34. The computer-system of claim 33, wherein the selected one or more chunks are distributed across two or more locations, and wherein each location corresponds to one of:

- a storage device added to a server from the set of servers used by the database system;
- a storage device belonging to a new server, wherein the new server is added to the set of servers used by the database system; or
- a network attached storage device of a remote server made accessible to a server from the set of servers used by the database system.

35. The computer-system of claim 33, wherein instructions for selecting one or more chunks according to criteria specified in the policy configuration comprise instructions for:

- for each of the one or more chunks, comparing the time attribute of the chunk with a threshold value.

40

36. The computer-system of claim 33, wherein the criteria specified in the policy configuration comprises age of the chunk, wherein age of a chunk is determined based on the set of values associated with the time attribute of the chunk.

37. The computer-system of claim 33, wherein the action associated with a selected chunk comprises deleting the selected chunk.

38. The computer-system of claim 33, wherein the action associated with a selected chunk comprises performing one or more of:

- adding an index to the selected chunk,
- dropping an index from the selected chunk,
- reindexing the selected chunk, or
- adding a constraint on the chunk.

39. The computer-system of claim 33, wherein the action associated with a selected chunk comprises moving the selected chunk from a first location to a second location.

40. The computer-system of claim 39, wherein the second location is determined based on the policy or the configuration of the selected chunk.

41. The computer-system of claim 39, wherein the second location is determined based on one or more of: access time of location, storage capacity of location, cost of location, or performance of location.

42. The computer-system of claim 33, wherein the action associated with a selected chunk comprises moving one or more records of the chunk to another chunk and updating the configurations of the selected chunk and the other chunk based on the movement of the one or more records.

43. The computer-system of claim 33, wherein a selected chunk has a plurality of replicas, each of the plurality of replicas storing a copy of a particular record.

44. The computer-system of claim 43, wherein the action is performed on one or more replicas from the plurality of replicas.

45. The computer-system of claim 43, wherein each of the plurality of replicas is stored in a different location.

46. The computer-system of claim 43, wherein the action comprises moving a replica from a first location to a second location, where the second location has different characteristics than the first location.

47. The computer-system of claim 43, wherein the number of replicas of a chunk are determined based on the configuration of the chunk.

48. The computer-system of claim 43, wherein the number of replicas of the chunk is determined based on a set of values associated with the time attribute for the chunk.

\* \* \* \* \*